

3D Virtual Worlds and the Metaverse: Current Status and Future Possibilities

JOHN DAVID N. DIONISIO

Loyola Marymount University

and

WILLIAM BURNS

Object Interoperability Lead for IEEE Virtual World Standard Group

and

RICHARD GILBERT

Loyola Marymount University

Moving from a set of independent virtual worlds to an integrated network of 3D virtual worlds or Metaverse rests on progress in four areas: immersive realism, ubiquity of access and identity, interoperability, and scalability. For each area, the current status and needed developments to achieve a functional Metaverse are described. Factors that support the formation of a viable Metaverse, such as institutional and popular interest and on-going improvements in hardware performance, and factors that constrain the achievement of this goal, including limits in computational methods and unrealized collaboration among virtual world stakeholders and developers, are also considered.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed applications*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Graphical user interfaces (GUI)*; *Voice I/O*; I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Virtual reality*; *Animation*

General Terms: Algorithms, Design, Human factors, Performance, Standardization

Additional Key Words and Phrases: Internet, immersion, virtual worlds, Second Life

1. INTRODUCTION TO VIRTUAL WORLDS AND THE METAVERSE

1.1 Definition and Significance of Virtual Worlds

Virtual worlds are persistent online computer-generated environments where multiple users in remote physical locations can interact in real time for the purposes of work or play. Virtual worlds constitute a subset of *virtual reality* applications, a more general term that refers to computer-generated simulations of three-dimensional objects or environments with seemingly real, direct, or physical user interaction. In 2008, the National Academy of Engineering (NAE) identified virtual reality as one of 14 Grand Challenges awaiting solutions in the 21st Century [Na-

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2011 ACM 0360-0300/11/1200-0001 \$5.00

tional Academy of Engineering 2008]. While the NAE Grand Challenges Committee has a predominantly American perspective, it is still informed by constituents with international backgrounds, with a UK member and several US-based members who have close associations with Latin American, African, and Asian institutions (<http://www.engineeringchallenges.org/cms/committee.aspx>). In addition, as [Zhao 2011] notes, a 2006 report by the Chinese government entitled *Development Plan Outline for Medium and Long-Term Science and Technology Development (2006-2020)* (http://www.gov.cn.jrzg/2006-02-09/content_183787.htm) and a 2007 Japanese government report labeled *Innovation 2025* (<http://www.cao.go.jp/innovation.index.html>) both included virtual reality as a priority technology worthy of development. Thus, there is international support for promoting advances in virtual reality applications.

1.2 Purpose and Scope of the Current Paper

This paper surveys the current status of computing as it applies to 3D virtual spaces and outlines what is needed to move from a set of independent virtual worlds to an integrated network of 3D virtual worlds or “Metaverse” that constitutes a compelling alternative realm for human socio-cultural interaction. In presenting this status report and roadmap for advancement, attention will be specifically directed to the following four features that are considered central components of a viable Metaverse:

- (1) *Realism*. Is the virtual space sufficiently realistic to enable users to feel psychologically and emotionally immersed in the alternative realm?
- (2) *Ubiquity*. Are the virtual spaces that comprise the Metaverse accessible through all existing digital devices (from desktops to tablets to mobile devices) and do the user’s virtual identities or collective persona remain intact throughout transitions within the Metaverse?
- (3) *Interoperability*. Do the virtual spaces employ standards such that (a) digital assets used in the reconstruction or rendering of virtual environments remain interchangeable across specific implementations and (b) users can move seamlessly between locations without interruption in their immersive experience?
- (4) *Scalability*. Does the server architecture deliver sufficient power to enable massive numbers of users to occupy the Metaverse without compromising the efficiency of the system and the experience of the users?

In order to provide context for considering the present state and potential future of 3D virtual spaces, the paper begins by presenting the historical development of virtual worlds and conceptions of the Metaverse. This history incorporates literary and gaming precursors to virtual world development, as well as direct advances in virtual world technology, because these literary and gaming developments often preceded, and significantly influenced, later technological achievements in virtual world technology. Thus, they are most accurately treated as important elements in the technical development of 3D spaces rather than as unrelated cultural events.

2. THE EVOLUTION OF VIRTUAL WORLDS

2.1 Five Phases of Virtual World Development

The development of virtual worlds has a detailed history in which acts of literary imagination and gaming innovations have led to advances in open-ended, socially oriented virtual platforms (see, for example, [Au 2008; Boellstorff 2008; Ludlow and Wallace 2007]). This history can be broadly divided into five phases.

In the initial phase, beginning in the late 1970s, text-based virtual worlds emerged in two varieties. MUDs or *multi-user dungeons* involved the creation of fantastic realities that resembled Tolkien’s *Lord of the Rings* or the role-playing dice game *Dungeons and Dragons*. MUSHs or *multi-user shared hallucinations* consisted of less defined, exploratory spaces where educators experimented with the possibility of collaborative creation [Turkle 1995].

The next phase of development occurred around a decade later when Lucasfilm, partly inspired by the 1984 publication of William Gibson’s *Neuromancer*, introduced Habitat for the Commodore 64 in 1986 and the Fujitsu platform in 1989. Habitat was the first high-profile commercial application of virtual world technology and the first virtual world to incorporate a graphical interface. This initial graphical interface was in 2D versus 3D, however, and the online environment appeared to the user as something analogous to a primitive cartoon operating at slow speeds via dial-up modems. Habitat was also the first virtual world to employ the term “avatar” (from the Sanskrit meaning “the deliberate appearance or manifestation of a deity on earth”) to describe its digital residents or inhabitants. As with the original Sanskrit, the contemporary usage of the term avatar involves a transposition of consciousness into a new form. However, in contrast to the ancient usage, the modern transition in form involves movement from a human body to a digital representation rather than from a god to a man or woman. Pioneering work in virtual reality systems and user interfaces mark the transition from this phase of development to the next [Blanchard et al. 1990; Cruz-Neira et al. 1992; Lanier 1992; Krueger 1993].

The third phase of development, beginning in the mid-1990s, was a particularly vibrant one with advances in computing power and graphics underlying progress in several major areas, including the introduction of user-created content, 3D graphics, open-ended socialization, and integrated audio. In 1994, Web World offered a 2.5D (isometric) world that provided users with open-ended building capability for the first time. The incorporation of user-based content creation tools ushered in a paradigm shift from pre-created virtual settings to online environments that were contributed to, changed, and built by participants in real time. In 1995, Worlds, Inc. (<http://www.worlds.com>) became the first publicly available virtual world with full three-dimensional graphics. Worlds, Inc. also revived the open-ended, non-game based genre (originally found in text-based MUSHs) by enabling users to socialize in 3D spaces, thus moving virtual worlds further away from a gaming model toward an emphasis on providing an alternative setting or culture to express the full range and complexity of human behavior. In this way, the scope and diversity of activities in virtual worlds came to parallel that of the Internet as a whole, only in 3D and with corresponding modes of interaction. 1995 also saw the introduction of Activeworlds (<http://www.activeworlds.com>), a virtual world

based entirely on the vision expressed in Neil Stephenson’s 1992 novel *Snow Crash*. Activeworlds explicitly expected users to personalize and co-construct a full 3D virtual environment. Finally, in 1996, OnLive! Traveler became the first publicly available 3D virtual environment to include natively utilized spatial voice chat and movement of avatar lips via processing phonemes.

The fourth phase of development occurred during the post-millennial decade. This period was characterized by dramatic expansion in the user-base of commercial virtual worlds such as Second Life, enhanced in-world content creation tools, increased involvement of major institutions from the physical world (e.g., corporations, colleges and universities, and non-profit organizations), the development of an advanced virtual economy, and gradual improvements in graphical fidelity. Avatar Reality’s Blue Mars, released in 2009, involved the most ambitious attempt to incorporate a much higher level of graphical realism into virtual worlds by using the then state-of-the-art CryEngine 2, which was originally developed by Crytek for gaming applications. Avatar Reality also allowed 3D objects to be created out-of-world as long as they could be saved in one of several open formats. Unfortunately, the effort to achieve greater graphical realism significantly raised the system requirements for client machines to a degree that was not cost-effective for Blue Mars’s target user base, and this overreaching necessitated massive restructuring/scaling down of a once promising initiative.

Overlapping the introduction of Second Life and Blue Mars has been a fifth phase of development. This phase, which began in 2007 and remains on-going, involves open source, decentralized contributions to the development of 3D virtual worlds. Solipsis is notable not only as one of the first open source virtual world systems but also for its proposed peer-to-peer architecture [Keller and Simon 2002; Frey et al. 2008]. It has since been followed by a host of other open source projects, including Open Cobalt [Open Cobalt Project 2011], Open Wonderland [Open Wonderland Foundation 2011], and OpenSimulator [OpenSimulator Project 2011].

Decentralized development has led to the decoupling of the client and server sides of a virtual world system, facilitated by convergence on the network protocol used by Linden Lab for Second Life as a de facto standard. The open source Imprudence/Kokua viewer marked the first “third-party viewer” for Second Life [Antonelli (AV) and Maxsted (AV) 2011]¹, an application category that grew to include projects such as Phoenix (and later Firestorm) [Lyon (AV) 2011], Singularity [Gearz (AV) 2011], Kirstens [Cinquetti (AV) 2011], and Nirans (itself derived from Kirstens) [Dean (AV) 2011]. On the server side, the aforementioned OpenSimulator has emerged as a Second Life workalike and alternative, and has itself spawned forks such as Aurora-Sim [Aurora-Sim Project 2011] and realXtend [realXtend Association 2011]. The Second Life protocol itself remains proprietary, under the full control of Linden Lab.

The emergent multiplicity of interoperable clients (viewers) and servers forms the

¹When individuals use an avatar name, as opposed to their physical world name, for a paper or software product, we have added a parenthetical (AV), for avatar, after that name to signify this choice. This conforms to the real world convention of citing an author’s chosen pseudonym (e.g., Mark Twain), rather than his or her given name (e.g., Samuel Clemens). Any author name that is not followed by an (AV) references the author’s physical world name.

basis of our fifth phase’s ultimate endpoint: full interoperability and interchangeability across virtual world servers and clients, in the same way that the Worldwide Web consists of multiple client (browser) and server options centered around the standard HTTP[S] protocol. Open source availability has facilitated new integration possibilities such as cloud-computing virtual world hosts and authentication using social network credentials [Trattner et al. 2010; Korolov 2011]. With activities such as this and continuing work on viewers and servers, it can be said that we have entered an open development phase for virtual worlds. Per Eric S. Raymond’s well-known metaphor, the virtual worlds cathedral has evolved into a virtual worlds bazaar [Raymond 2001].

Tables I to III summarize the major literary/narrative influences to virtual world development and the advances in virtual world technology that have occurred during the 30 plus years since MUDs and MUSHs were introduced.

YEAR(S)	LITERARY/NARRATIVE	SIGNIFICANCE
1955/1966	<i>Lord of the Rings</i>	Tolkien’s foundational work of 20th-century fantasy literature serves as a source of inspiration for many games and virtual world environments.
1974	<i>Dungeons and Dragons</i>	Originally designed by Gary Gygax and Dave Arneson. Widely regarded as the beginning of modern role-playing games.
1981	<i>True Names</i>	Vernor Vinge’s science fiction novella presents a fully fleshed-out version of cyberspace. Influences later classics such as <i>Neuromancer</i> and <i>Snow Crash</i> .
1984	<i>Neuromancer</i>	William Gibson’s seminal cyberpunk novel popularized the early concept of cyberspace as “the Matrix.”
1992	<i>Snow Crash</i>	The term <i>Metaverse</i> was coined in Neal Stephenson’s science fiction novel to describe a virtual reality-based successor to the Internet.

Table I. Major literary/narrative influences to virtual world technology

Reflecting these developments, Gilbert has identified 5 essential features that characterize a contemporary, state-of-the art, virtual world [Gilbert 2011]:

- (1) It has a *3D graphical interface* and *integrated audio*: An environment with a text interface alone does not constitute an advanced virtual world.
- (2) It supports *massively multi-user remote interactivity*: Simultaneous interactivity among large numbers of users in remote physical locations is a minimum requirement, not an advanced feature.
- (3) It is *persistent*: The virtual environment continues to operate even when a particular user is not connected.
- (4) It is *immersive*: The environment’s level of spatial, environmental, and multi-sensory realism creates a sense of psychological presence. Users have a sense of “being inside,” “inhabiting,” or “residing within” the digital environment rather than being outside of it, thus intensifying their psychological experience.

Phase I: Text-based virtual worlds — Late 1970s		
YEARS	EVENT/MILESTONE	SIGNIFICANCE
1979	<i>MUDs and MUSHes</i>	Roy Trubshaw and Richard Bartle complete the first “multi-user dungeon” or “multi-user domain” — a multiplayer, real-time, virtual gaming world described primarily in text. MUSHes, text-based on-line environments in which multiple users are connected at the same time for open-ended socialization and collaborative work, also appear at this time.
Phase II: Graphical interface and commercial application — 1980s		
1986/1989	<i>Habitat for Commodore 64 (1986) and Fujitsu platform (1989)</i>	Partly inspired by William Gibson’s <i>Neuromancer</i> . The first commercial simulated multi-user environment using 2D graphical representations and employing the term “avatar,” borrowed from the Sanskrit term meaning the deliberate appearance or manifestation of a deity in human form.
early 1990s	<i>Reality Built For Two, CAVE, Artificial Reality</i>	Prototype virtual reality systems and immersive environments begin to proliferate.
Phase III: User-created content, 3D graphics, open-ended socialization, integrated audio — 1990s		
1994	<i>Web World</i>	The first 2.5D (isometric) world where tens of thousands could chat, build, and travel. Initiated a paradigm shift from pre-created environments to environments contributed to, changed, and built by participants in real time.
1995	<i>Worlds Inc.</i>	One of the first publicly available 3D virtual user environments. Enhanced the open-ended, non-game-based genre by enabling users to socialize in 3D spaces. Continued moving virtual worlds away from a gaming model toward an emphasis on providing an alternative setting or culture to express the full range and complexity of human behavior.
1995	<i>Activeworlds</i>	Based entirely on <i>Snow Crash</i> , popularized the project of creating an actual Metaverse. Offered basic content-creation tools to personalize and co-construct the virtual environment.
1996	<i>OnLive! Traveler</i>	The first publicly available system that natively utilized spatial voice chat and incorporated the movement of avatar lips via processing phonemes.

Table II. Major 20th-century advances in virtual world technology

- (5) It emphasizes *user-generated activities and goals* and provides *content creation tools* for personalization of the virtual environment and experience. In contrast to immersive games, where goals — such as amassing points, overcoming an enemy, or fulfilling a quest — are built into the program, virtual worlds provide a more open-ended setting where, similar to physical life and culture, users can define and implement their own activities and goals. However, there are some

Phase IV: Major expansion in commercial virtual world user bases, enhanced content creation tools, increased institutional presence, development of robust virtual economy, improvements in graphical fidelity — Post-millennial decade

YEARS	EVENT/MILESTONE	SIGNIFICANCE
2003–present	<i>Second Life</i>	Popular open-ended, commercial virtual environment with (1) in-world live editing, (2) ability to import externally created 3D objects into the virtual environment, and (3) advanced virtual economy. Primary virtual world for corporate and educational institutions.
2009–present	<i>Avatar Reality/ Blue Mars</i>	A closed-source foray into much higher graphical realism using 3D graphics engine technology initially developed in the gaming industry.

Phase V: Open, decentralized development — 2007 and beyond

2007	<i>Solipsis</i>	The first open source decentralized virtual world system. Open source development theoretically allowed new server and client variants to be created from the original code base — something that would actually become a reality with systems other than Solipsis.
2008	<i>Imprudence/Kokua</i>	One of the earliest alternative, open source viewers for an existing virtual world server (Second Life). So-called “third-party viewers” for Second Life have proliferated ever since.
2009	<i>OpenSimulator</i>	First instance of multiple servers following the same virtual world protocol (Second Life), later accompanied by a choice of multiple viewers that use this protocol. While the Second Life protocol has become a de facto standard, the protocol itself remains proprietary and occasionally requires reverse engineering.
2010 and beyond	<i>Open Development of the Metaverse</i>	Interoperability and interchangeability across servers and clients through standard virtual world protocols, formats, and digital credentials — users and providers can choose among virtual world client and server implementations, respectively, without worrying about compatibility or authentication, just as today’s web runs multiple web servers and browsers with standards such as OpenID and OAuth.

Table III. Major 21st-century advances in virtual world technology

immersive games, such as World of Warcraft, that include specific goals for the users but are still psychologically and socially complex. These intricate immersive environments stand at the border of games and virtual worlds.

2.2 From Individual Virtual Worlds to the Metaverse

The word Metaverse is a portmanteau of the prefix “meta” (meaning “beyond”) and the suffix “verse” (shorthand for “universe”). Thus, it literally means a universe beyond the physical world. More specifically, this “universe beyond” refers to a computer-generated world, distinguishing it from metaphysical or spiritual conceptions of domains beyond the physical realm. In addition, the Metaverse refers

to a fully immersive, three-dimensional digital environment in contrast to the more inclusive concept of “cyberspace” that reflects the totality of shared online space across all dimensions of representation.

While the Metaverse always references an immersive, three-dimensional, digital space, conceptions about its specific nature and organization have changed over time. The general progression has been from viewing the Metaverse as an amplified version of an individual virtual world to conceiving it as a large network of interconnected virtual worlds. Neal Stephenson, who coined the term in his 1992 novel, *Snow Crash*, vividly conveyed the *Metaverse as Virtual World* perspective. In Stephenson’s conception of the Metaverse, humans-as-avatars interact with intelligent agents and each other in an immersive world that appears as a nighttime metropolis developed along a neon-lit, hundred-meter wide, grand boulevard called the Street, evoking images of an exaggerated Las Vegas strip. The Street runs the entire circumference of a featureless, black planet considerably larger than Earth that has been visited by 120 million users, approximately 15 million of whom occupy the Street at a given time. Users gain access to the Metaverse through computer terminals that project a first person perspective virtual reality display onto goggles and pump stereo digital sound into small earphones that drop from the bows of the goggles and plug into the user’s ears. Users have the ability to customize their avatars with the sole restriction of height (to avoid mile-high avatars), to travel by walking or by virtual vehicle, to build structures on acquired parcels of virtual real estate, and to engage in the full range of human social and instrumental activities. Thus, the Metaverse that Stephenson brilliantly imagined is, in both form and operation, essentially an extremely large and heavily populated virtual world that operates, not as a gaming environment with specific parameters and goals, but as an open-ended digital culture that operates in parallel with the physical domain.

Since Stephenson’s novel appeared, technological advances have enabled real-life implementation of virtual worlds and more complex and expansive conceptions of the Metaverse have developed. In 2007, The Metaverse Roadmap Project [Smart et al. 2007] offered a multi-faceted conception of the Metaverse that involved both “simulation technologies that create physically persistent virtual spaces such as virtual and mirror worlds” and “technologies that virtually-enhance physical reality such as augmented reality” (i.e., technologies that connect networked information and computational intelligence to physical objects and spaces). While this effort is notable in its attempt to view the Metaverse in broader terms than an individual virtual world, the inclusion of augmented reality technologies served to redirect attention from the core qualities of immersion, three-dimensionality, and simulation that are the foundations of virtual world environments. We consider the augmented reality space to be a subset of the Metaverse that constitutes a crossroads between purely virtual environments and purely real or visceral environments. Like any virtual world system, augmented reality constructs also access assets and data from a self-contained or shared world state, overlaying them on a view of the physical world rather than a synthetic one.

In contrast to the Metaverse Roadmap, a 2008 white paper on Solipsis, an open source architecture for creating large systems of virtual environments using a peer-to-peer topology, provided the first published account of the contemporary *Meta-*ACM Computing Surveys, Vol. V, No. N, December 2011.

verse as Network of Virtual Worlds perspective. The Solipsis white paper defined the concept as “a massive infrastructure of inter-linked virtual worlds accessible via a common user interface (browser) and incorporating both 2D and 3D in an Immersive Internet” [Frey et al. 2008]. Frey et al. and the IEEE Virtual World Standard Group (<http://www.metaversrsestandards.org>) also offered a clear developmental progression from an individual virtual world to the Metaverse using concepts and terminology aligned with the organization of the physical universe [Burns 2010; IEEE VW Standard Working Group 2011b]. This progression starts with separate virtual worlds or “MetaWorlds” (analogous to individual physical planets) with no inter-world transit capabilities (e.g., Second Life, Entropia Universe, and the Chinese virtual world of Hipihi). “MetaGalaxies” (sometimes referred to as “hypergrids”) then involve multiple virtual worlds clustered together as perceived collectives under a single authority. MetaGalaxies such as Activeworlds and OpenSim Hypergrid-enabled virtual environments permit perceived teleportation or space travel. In such environments, users have the perception of leaving one “planet” or perceived contiguous space and arriving at another. The progression of development culminates in a complete Metaverse that involves multiple MetaGalaxies and MetaWorld systems. A standardized protocol and set of abilities would allow users to move between virtual worlds in a seamless manner regardless of the controlling entity for any particular virtual region.

2.3 Computational Advances and the Metaverse

While there are multiple virtual worlds currently in existence, and arguably one hypergrid or MetaGalaxy in Activeworlds, the Metaverse itself remains a concept awaiting full implementation. As noted in Section 1.2, the creation of a fully realized Metaverse will rest on continued progress with regard to four essential features of virtual world technology: psychological realism, ubiquity of access and identity, interoperability of content and experience across virtual environments, and scalability. In the section to follow, the current status of each of these features will be presented along with developments in each area that could contribute to the attainment of a viable, psychologically compelling Metaverse.

In considering this framework for advancing virtual world technology, it is important to note two qualifications. First, in the interest of clarity of organization and presentation, the four features of the Metaverse will be mainly presented as separate and distinct variables even though overlap exists among these factors (e.g., increased graphical realism will have implications for server requirements and scalability).

Second, it is important to acknowledge the likelihood that rapid changes in computing (as popularized by Moore’s law and Kurzweil’s law of accelerating returns [Moore 1965; Kurzweil 2001]) will impact the currency of some aspects of this survey. Thus, examples will be cited from time to time for illustrative purposes with full knowledge that the time it takes to prepare, review, and disseminate this document will render some of these examples outdated or obsolete. Nevertheless, while many specifics that make up the virtual realm are subject to rapid change, the general principles and computing issues related to the formation of a fully operational Metaverse that are the primary focus of the present work will continue to be relevant over time [Lanier and Biocca 1992; National Academy of Engineering

2008; Zhao 2011].

3. FEATURES OF THE METAVERSE: CURRENT STATUS AND FUTURE POSSIBILITIES

3.1 Realism

We immediately qualify our use of the term “realism” to mean, in this context, *immersive* realism. In the same way that “realism” in cinematic computer-generated imagery is qualified by its *believability* rather than devotion to detail (though certainly a sufficient degree of detail is expected), “realism” in the Metaverse is sought in the service of a user’s psychological and emotional engagement within the environment. A virtual environment is perceived as more realistic based on the degree to which it transports a user into that environment, and the transparency of the boundary between the user’s physical actions and those of his or her avatar. By this token, virtual world realism is not purely additive nor, visually speaking, strictly photographic; in many cases, strategic rendering choices can yield better returns than merely adding polygons, pixels, objects, or bits in general across the board.

What does remain constant across all perspectives on realism is the “instrumentation” through which human beings interact with the environment: their senses and their bodies, particularly through their faces and hands. We thus approach realism through this structure of perception and expression.

We treat the subject at a historical and survey level in this article, examining current and future trends. Greater technical detail and focus solely on the [perceived] realism of virtual worlds can be found in the cited work and tutorial materials such as [Glencross et al. 2006].

3.1.1 *Sight.* Not surprisingly, sight and visuals comprise the sensory channel with the most extensive history within virtual worlds. As described in Section 2, the earliest visual medium for virtual worlds, and for all of computing, for that matter, was plain text. Text was used to form imagery within the “mind’s eye,” and to a certain degree it remains a very effective mechanism for doing so. In the end, however, words and symbols are indirect: they *describe* a world, and leave specifics to individuals. Effective visual immersion involves eliminating this level of indirection — a virtual world’s visual presentation seeks to be as information-rich to our eyes as the real world is. The brain then recognizes imagery rather than synthesizing or recalling it (as it would when reading text).

To this end, the state of the art in visual immersion has, until recently, hewn very closely to the state of the art in real-time computer graphics. The precise meaning of real-time visual perception is complicated and fraught with nuances and external factors such as attention, fixation, and non-visual cues [Hayhoe et al. 2002; Recanzone 2003; Chow et al. 2006]. For this discussion, “real-time” refers to the simplified metric of *frame rate*, or the frequency with which a computer graphics system can render a scene. As graphics hardware and algorithms have pushed the threshold of what can be computed and displayed at 30 or more frames per second — a typically accepted minimum frame rate for real-time perception [Kumar et al. 2008; Liu et al. 2010] — virtual world applications have provided increasingly detailed visuals commensurate with the techniques of the time. Thus, alongside games, 3D modeling software, and, to a certain extent, 3D cinematic

animation (qualified as such because the real-time constraint is lifted for the final product), virtual worlds have seen a progression of visual detail from flat polygons to smooth shading and texture mapping, and finally to programmable shaders.

Initially, visual richness was accomplished simply by increasing the amount of data used to render a scene: more and smaller polygons, higher-resolution textures, and interpolation techniques such as anti-aliasing and smooth shading. The fixed-function graphics pipeline, so named because it processed 3D models in a uniform manner, facilitated the stabilization of graphics libraries and incremental performance improvements in graphics hardware, but it also restricted the options available for delivering detailed objects and realistic rendering [Olano et al. 2004; Olano and Lastra 1998]. This served the area well for a time, as hardware improvements accommodated the increasing bandwidth and standards emerged for representing and storing 3D models. Ultimately, however, using increased data for increased detail resulted in diminishing returns, as the cost of developing the information (3D models, 2D textures, animation “rigging”) can easily surpass the benefit seen in visual realism.

The arrival of programmable shaders, and their eventual establishment as a new functional baseline for 3D libraries such as OpenGL, represented a significant step for virtual world visuals, as it did for computer graphics applications in general [Rost et al. 2009]. With programmable shaders, many aspects of 3D models, whether in terms of their geometry (vertex shaders) or their final image rendering (fragment shaders), became expressible algorithmically, thus eliminating the need for larger numbers of polygons or textures, while at the same time sharply expanding the variety and flexibility with which objects can be rendered or presented. This degree of flexibility has produced compelling real-time techniques for a variety of object and scene types. The separation of vertex and fragment shaders has also facilitated the augmentation of traditional approaches involving geometry and object models with image-space techniques [Roh and Lee 2005; Shah 2007].

Some of these techniques can see immediate applicability in virtual world environments, as they involve exterior environments such as terrain, lighting, bodies of water, precipitation, forests, fields, and the atmosphere [Hu et al. 2006; Wang et al. 2006; Bouthors et al. 2008; Boulanger et al. 2009; Seng and Wang 2009; Elek and Kmoch 2010] or common objects, materials, or phenomena such as clothing, reflected light, smoke, translucence, or gloss [Adabala et al. 2003; Lewinski 2011; Sun and Mukherjee 2006; Ghosh 2007; Shah et al. 2009]. Techniques for specialized applications such as surgery simulation [Hao et al. 2009] and radioactive threat training [Koepnick et al. 2010] have a place as well, since virtual world environments strive to be as open-ended as possible, with minimal restrictions on the types of activities that can be performed within them (just as in the real world).

Most virtual world client applications or *viewers* currently stand in transition between fixed function rendering (explicit polygon representations, texture maps for small details) and programmable shaders (particle generation, bump mapping, atmosphere and lighting, water) [Linden Lab 2011a]. Virtual world viewers tend to lag behind other graphics applications such as games, visual effects, or 3D modeling. Real-time lighting and shadows, for instance, were not implemented in the official Second Life viewer until June 2011, long after these had become fairly standard in

other graphics application categories [Linden Lab 2011b].

The relative dominance of Second Life as a virtual world platform has resulted in an “ecology” of viewers that can all connect to the same Second Life environment but do so with varying features and functionalities, some of which pertain to visual detail [Antonelli (AV) and Maxsted (AV) 2011; Gearz (AV) 2011; Lyon (AV) 2011].² In particular, Kirstens Viewer [Cinquetti (AV) 2011] has, as one of its explicit goals, the extension and integration of graphics techniques such as realistic shadows, depth of field, and other functionalities that take better advantage of the flexibility afforded by programmable shaders.

It was stated earlier that visual immersion has hewn very closely to real-time computer graphics *until recently*. This divergence corresponds to the aforementioned “type” of realism that is demanded by virtual world applications, which is not necessarily detail, but more particularly a sense of transference, of transportation into a different environment or perceived reality. It can be stated that the visual fidelity afforded by improved processing power, display hardware, and graphics algorithms (particularly as enabled by programmable shaders) has crossed a certain threshold of sufficiency, such that improving the sense of visual immersion has shifted from pure detail to specific visual elements.

One might consider, then, that stereoscopic vision may be a key component of visual immersion. This, to a degree, has been borne out by the recent surge in 3D films, and particularly exemplified by the success of films with genuine three-dimensional data, such as *Avatar* and computer-animated work.

A limiting factor of 3D viewing, however, is the need for special glasses in order to fully isolate the imagery seen by the left and right eyes. While such glasses have become progressively less obtrusive, and may eventually become completely unnecessary, they still represent a hindrance of sorts, and represent a variant of the “come as you are” constraint taken from gesture-based interfaces, which states that, ideally, user interfaces must minimize the need for attachments or special constraints [Triesch and von der Malsburg 1998; Wachs et al. 2011].

3.1.2 Sound. Like visual realism, the digital replication or generation of sound finds applications well beyond those of virtual worlds, and can be argued to have had a greater impact on society as a whole than computer-generated visuals, thanks to the CD player, MP3 and its successors, and home theater systems. Unlike visuals, which are, quite literally, “in one’s face” when interacting with current virtual world environments, the role of audio is quite dual in nature, consisting of distinctly conscious and un- or subconscious sources of sound:

- (1) “Conscious” or “front-and-center” audio in virtual worlds consists primarily of speech: speaking and listening are our most natural and facile forms of verbal communication, and the availability of speech is a significant factor in psychologically immersing and engaging ourselves in a virtual world, since it engages us directly with other inhabitants of that environment, perhaps more so than reading their (virtual) faces, posture, and movement.
- (2) “Ambient” audio in virtual worlds consists of sound that we do not consciously

²This multiplicity of viewers for the same information space is analogous to the availability of multiple browser options for the Worldwide Web.

process, but whose presence or absence subtly influences the sense of immersion within that environment. This sense of sound immersion derives directly from the aural environment of the real world: we are, at all times, enveloped in sound, whether or not we are consciously listening to it. Such sound provides important positional and spatial cues, the perception of which contributes significantly to our sense of “placement” within a particular situation [Blauert 1996].

For the first, “conscious” form of audio stimulus, we note that high-quality voice chat which reasonably captures the nuances of verbal communication is now generally available [Short et al. 1976; Lober et al. 2007; Wadley et al. 2009]. In this regard, the sense of hearing is somewhat ahead of sight in capturing avatar expression and interaction. Future avenues in this area include voice masking and modulation technologies that not only capture the user’s spoken expressions, but also customize the reproduced vocals seamlessly and effectively [Ueda et al. 2006; Xu et al. 2008; Mecwan et al. 2009].

Ambient audio, however, merits additional discussion. The accurate reproduction of *individual* sounds serves as a starting point: sound sample quality is limited today solely by storage and bandwidth availability, as well as the dynamic range of audio output devices. However, as anyone who has listened to an extremely faithful, lossless recording on a single speaker will attest, such one-dimensional accuracy is only the beginning.

In some respects, the *immersiveness* of sound can be seen as a broader-based pursuit than immersive visuals, with the ubiquity of sound systems on everything from mobile devices to game consoles to PCs to top-of-the-line studios and theaters. Over the years, “stereo” has given way to “2.1,” “5.1,” “7.1” systems and beyond, where the first value indicates the number of distinct directional channels and the second represents the number of omnidirectional, low-frequency channels. The very buzzword “surround sound” communicates our desire to be completely enveloped within an audio environment — and we know it when we hear it. We just don’t hear it very often, outside of the real world itself.

The correlation of “immersiveness” to “number of directional channels” is a straightforward, if ultimately simplistic, approximation of how we perceive sound. Sound from the environment reaches our two ears at subtly different levels and frequencies, differences which the human brain can seamlessly merge and parse into strikingly complete positional information [Blauert 1996]. Multichannel sound attempts to replicate these variations by literally projecting different levels of audio from different locations. However, genuine three-dimensional sound perception is more complicated than mere division into channels, as such perception is subject to highly individualized factors as the size and shape of the head, the shape of the ears, and years or decades of training and fine-tuning.

The study of *binaural sound*, as opposed to discrete multichannel reproduction, seeks to record and reproduce sound in a manner that is more faithful to the way we hear in three dimensions. Such precision is modeled as a *head-related transfer function* (HRTF) that specifies how the ear receives sound from a particular point in space. The HRTF captures the effect of the myriad variables involved in true spatial hearing — acoustic properties of the head, ears, etc. — and determines precisely what must reach the left and right ears of the listener. As a result,

binaural audio also requires headphones to have the desired effect, much the same way that stereoscopic vision requires glasses. Both sensory approaches require the precise isolation and coordination of stimuli delivered to the organ involved, so that the intended perceptual “computation” occurs in the brain.

The theory and technology of binaural audio has actually been around for a strikingly long time, with the first such recordings being made as far back as 1881 and occasional milestones and products being developed through the 20th century [Kall Binaural Audio 2010]. Not surprisingly, binaural audio synthesis, thanks to the mathematics of HRTFs, has also been studied and implemented [Brown and Duda 1998; Cobos et al. 2010; Action = Reaction Labs 2010]. What *is* surprising is its lack of application in virtual world technology, with no known systems adopting binaural synthesis techniques to create a genuinely spatial sound field for its users. The implementation of binaural audio in virtual worlds thus represents a clear signpost for increased realism (and thus immersiveness) in this area.

Like 3D viewing, binaural audio requires clean isolation of left- and right-side signals in order to be perceived properly, with headphones serving as the aural analog of 3D glasses. Unlike 3D viewing, individual physical differences affect the perceived sound field as well, such that perfect binaural listening cannot really be achieved unless recording conditions (or synthesis parameters) accurately model the listener’s anatomy and environment. Algorithmic tuning within the spatial audio environment in order to allow each listener to specifically match his/her hearing situation would be ideal. Much like we train voice recognition, or use an equalizer in audio for best effect to the listener, binaural settings and tuning would allow the listener to adjust to their specific circumstances and gain the most from the experience. It can be argued, however, that the increased spatial fidelity that is achieved even by non-ideal binaural sound would still represent appreciable progress in the audio realism of a virtual world application.

Binaural sound synthesis algorithms take, as input, a pre-recorded sound sample, positional information, and one or more HRTFs. Multiplexing these samples may then produce the overall spatial audio environment for a virtual world. However, in much the same way that image textures of static resolution and variety quickly reveal their limitations as tiles and blocks, so does a collection of mixed-and-matched samples eventually exhibit a degree of repetitiveness, even when presented spatially.

An additional layer of audio realism can be achieved through synthesis of the actual *samples* based on the characteristics of the objects generating the sound. Wind blowing through a tunnel will sound different from wind blowing through a field; a footfall should vary based on the footwear and the ground being traversed. Though subtle, such audio cues influence our perception of an environment’s realism, and thus our sense of immersion within it. When placed in the context of an open-ended virtual world, does this obligate the system to load up a vast library of pre-recorded samples? In the same way that modern visuals cannot rely solely on pre-drawn textures for detail, one can say that the ambient sound environment should also explore the dynamic synthesis of sound effects — a computational foley artist, so to speak [van den Doel et al. 2001].

Work on precisely this type of dynamic sound synthesis, based on factors ranging from rigid body interactions such as collisions, fractures, scratching, or sliding

	<i>Sight</i>	<i>Sound</i>
<i>Static representation</i>	textures, fixed function	pre-recorded samples, discrete channels
<i>Procedural approach</i>	programmable shaders	dynamic sound synthesis
<i>Spatial presentation</i>	stereoscopic vision	binaural audio
<i>Spatial equipment</i>	goggles/glasses	headphones
<i>Avatar expression</i>	face, gesture	voice chat, masking

Table IV. Conceptual equivalents in visual and audio realism.

[O’Brien et al. 2002; Bonneel et al. 2008; Chadwick et al. 2009; Zheng and James 2010] to adjustments based on the rendered texture of an object [Schmidt and Dionisio 2008], exists in the literature. Synthesis of fluid sounds is also of particular interest, as bodies of water and other liquids are sometimes prominent components of a virtual environment [Zheng and James 2009; Moss et al. 2010]. These techniques have not, however, seen much implementation in the virtual world milieu.

Imagine the full effect of a virtual world audio environment where every subtle interaction of avatars and objects produces a reasonably unique and believable sound, which is then rendered, through real-time sound propagation algorithms and binaural synthesis, in a manner that accurately reflects its spatial relationship with the user’s avatar. This research direction for realistic sound within the Metaverse has begun to produce results, and constitutes a desirable endpoint for truly immersive audio [Taylor et al. 2009].

Table IV seeks to summarize the realism discussion thus far, highlighting technologies and approaches that are analogous to each other across these senses.

3.1.3 Touch. After sight and sound, the sense of touch gets most of the remaining attention in terms of virtual world realism. It turns out that, in virtual worlds, the very notion of “touch” is more open to interpretation than sight and sound.

The classic, and perhaps most literal, notion of touch in virtual environments is *haptic* or *force feedback*. The idea behind haptic feedback is to convert virtual contacts into physical ones. Force feedback is a specific, somewhat simpler form of haptic feedback that pertains specifically to having physical devices push against or resist the user’s body, primarily his or her hands or arms. *Tactile feedback* has been differentiated from haptic/force feedback as involving only the skin as opposed to the muscles, joints, and entire body parts [Subramanian et al. 2005].

Haptics are a tale of two technologies: on the one hand, a very simple form of haptic feedback is widely available in game consoles and controllers, in the form of vibrations timed to specific events within a game or virtual environment. In addition, user input delivered by movement and gestures, which by nature produces a sort of “collateral” haptic feedback, has also become fairly common, available in two forms: hand-in-hand with the vibrating controls of the same consoles and systems, and alongside mobile and ultramobile devices that are equipped with gyroscopes and accelerometers. Microsoft’s Project Natal, which eventually became the Kinect product, took movement interfaces to a new phase with the first mainstream “come-as-you-are” motion system: users no longer require a handheld or attached controller in order to convey movement.

While this form of haptic feedback has become fairly commonplace, in many respects such sensations remain limited, synthetic, and ultimately not very immer-

sive. Feeling vibrations when colliding with objects or firing weapons, for example, constitutes a form of sensory mismatch, and is ultimately interpreted as an indirect cue — hardly more realistic than, say, a system beep or an icon on the screen. Movement interfaces, while providing some degree of inertial or muscular sensation, ultimately only go halfway: if a sweeping hand gesture corresponds to a push in the virtual environment, the *reaction* to the user’s *action* remains relegated to visual or aural cues, or, at most, a vibration on a handheld controller.

Thus, targeted haptic feedback — bodily sensations that map more naturally to the virtual environment — remains an area of research and investigation. The work ranges from implementations for specific applications [Harders et al. 2007; Abate et al. 2009] to work that is more integrated with existing platforms, facilitated by open source development [de Pascale et al. 2008]. Other work focuses on the haptic devices themselves — that is, the apparatus for *delivering* force feedback to the user [Folgheraiter et al. 2008; Bullion and Gurocak 2009; Withana et al. 2010]. Finally, the very methodology of capturing, modeling, and finally reproducing touch-related object properties has been proposed as *haptography* [Kuchenbecker 2008].

Any discussion of haptic or tactile feedback in virtual worlds would be incomplete without considering touch stimuli of a sexual nature. Sexually-motivated applications have been an “open secret” of technology adoption, and virtual worlds are no exception [Gilbert et al. 2011]. So-called “teledildonic” devices such as the Simulator and the Fleshlight seek to convey tactile feedback to the genital area, with the goal of promoting the realism of virtual world sexual encounters [Lynn 2004]. Unfortunately, empirical studies of their effectiveness and overall role in virtual environments, if any, remain unpublished or unavailable.

Less sexual, but potentially just as intimate, are avatar interactions such as hugs, cuddles, and tickles. Explorations in this direction have been made as well, in an area termed *mediated social touch* [Haans and IJsselsteijn 2006]. Such interactions require *haptic garments* or *jackets* that provide force or tactile feedback beyond the user’s hands [Tsetsrukou 2010; Rahman et al. 2010]. A “suite” of such devices has even been integrated with Linden Lab’s open source Second Life viewer to facilitate multimodal, affective, non-verbal communication among avatars in that virtual world [Tsetsrukou et al. 2010]. While such technologies still require additional equipment that may limit widespread adoption and use, they do reflect the aforementioned focus on enhanced avatar interaction as a particular component of virtual world realism and immersion.

Despite the work in this area, the very presence of a physical device may ultimately dilute users’ sense of immersion (another manifestation of the “come as you are” principle). Even with such attachments, no single device can yet convey the full spectrum of force and tactile stimuli that a real-world environment conveys. As a result, some studies have shown that user acceptance of such devices remains somewhat lukewarm or indifferent [Krol et al. 2009]. For networked virtual worlds in particular, technical issues such as latency detract from the effectiveness of haptic feedback [Jay et al. 2007]. The work continues, however, as other studies have noted that haptic feedback, once it reaches sufficient fidelity to real-life interactions, does indeed increase one’s feeling of social presence [Sallnäs 2010].

The physical constraints of “true” haptic feedback, particularly as imposed by

the need for physical attachments to the user, has triggered some interesting alternative paths. *Pseudo-haptic feedback* seeks to convey the *perception* of haptics by exploiting connections between visual and haptic stimuli [Lécuyer 2009]. Further, the responsiveness of objects to an avatar’s proximity and collisions conveys a sense of involvement and interaction that may stand in for actual haptic or force feedback. Such *proxemic interactions* are in fact sufficiently compelling that research in this area is being conducted even for real world systems [Greenberg et al. 2011]. Interestingly, *virtual proxemic interfaces* have an advantage of sorts over their real-world counterparts: in a virtual environment, sensing the state of the world is a non-issue, since the environment itself manages that state. Having such an accurate reading of virtual objects’ and avatars’ locations, velocities, materials, and shapes can facilitate very rich and detailed proxemic behavior. Such behavior can in turn promote a sense of touch-related immersion without having to deliver actual force or touch stimuli at all.

3.1.4 *Other Senses and Stimuli.* Smell and taste, as well as other perceived stimuli such as balance, acceleration, temperature, kinaesthesia, direction, time, and others, have not generally been as well-served by virtual world technologies as sight, sound, and touch. Of these, balance and acceleration have perhaps seen the most attention, through motion and flight simulators. Such simulators take “immersion” literally, with complete physical enclosure — ultimately not practical for all but the most devout (and wealthy!) of virtual world users.

3.1.5 *Gestures and Expressions.* It can be argued that the presence of and interactions with fellow human beings, *as mediated through virtual world avatars*, contributes strongly to the immersive experience that is deemed central to virtual world realism. As such, the more natural and expressive an avatar seems to be, the greater the perceived “reality” of the virtual environment. Traditional rendered detail for avatars has given way to affective cues like poses, incidental sounds, and facial expressions. Such cues range from the subtle, such as blinking, to the integrative, such as displaying mouth movement while audio chat is taking place. Anything that can make the avatar look more “alive” is worth exploring, as the appearance of life (or liveliness) translates to the sense of immersion [Paleari and Lisetti 2006; Martin et al. 2007; Geddes 2010]. Blascovich and Bailenson use the term *human social agency* to denote this characteristic. The primary variables of such agency have been stratified as *movement realism* (gestures, expressions, postures, etc.), *anthropometric realism* (recognizable human body parts), and *photographic realism* (closeness to actual human appearance), with the first two variables carrying greater weight than the third [Blascovich and Bailenson 2011].

Based on this premise, emerging work on the real-time rendering of human beings is anticipated as the next great sensory surge for virtual worlds. Such work includes increasingly realistic display and animation of the human form, not only of the body itself but of clothes and other attachments [Stoll et al. 2010; Lee et al. 2010]. In this effort, the face may be a particular area of focus — perhaps the new window to the “virtual soul,” at least until the rendering of eyes alone can capture the richness and depth of human emotion [Chun et al. 2007; Jimenez et al. 2010].

Alongside the *output* of improved visuals, audio, and touch is the requirement for

the non-intrusive *input* of the data that inform such sensory stimuli, particularly with respect to avatar interaction and communication. As mentioned, many sensory technologies rely on additional equipment beyond standard computing devices, ranging from the relatively common (headphones) to the rare or inconvenient (haptic garments) — and these are just for sensory *output*. Input devices are currently even more limited (keyboard, mouse) or somewhat untapped (camera), with only the microphone, in the audio category, being close to fulfilling the multiple constraints that an input device be highly available and unobtrusive, while delivering relative fidelity and integration into a virtual world environment.

Multiple research opportunities thus exist for improving the sensory *input* that can be used by virtual worlds, especially the kinds of input that enhance interaction among avatars, which, as we have proposed, is a crucial contributor to a system’s immersiveness and thus realism. On the visual side, capturing the actual facial expression of the user at a given moment [Chandrasiri et al. 2004; Chun et al. 2007; Bradley et al. 2010] serves as an ideal complement to aforementioned facial rendering algorithms [Jimenez et al. 2010] and effectively removes artificial indirection or mediation via keyboard or mouse for facial expressions between real and virtual environments. Some work covers the whole cycle, again indicative of the synergy between capture and presentation [Sung and Kim 2009].

Some input approaches are multimodal, integrating input for one sense with output for another. For example, head tracking, motion detection, or full-body video capture can produce more natural avatar animations [Kawato and Ohya 2000; Newman et al. 2000; Kim et al. 2010] and more immersive or intuitive user interfaces [Francone and Nigay 2011; Rankine et al. 2011]. Such cues can also be synchronized with supporting stimuli such as sounds (speech, breathing), environmental visuals (depth of field, subjective lighting), and other aspects of the virtual environment [Arangarasan and Phillips Jr. 2002; Latoschik 2005]. Integrating technologies so that the whole becomes greater than the sum of its parts is a theme that is particularly compelling for virtual worlds — a supposition that not only makes intuitive sense but has also seen support from published studies [Biocca et al. 2001; Naumann et al. 2010] and integrative presentations on the subject at conferences, particularly at SIGGRAPH [Fisher et al. 2004; Otaduy et al. 2009].

The need to have multiple devices, wearables, and other equipment for delivering these stimuli diverges from the “come as you are” principle and poses a barrier to widespread use. This has motivated the pursuit of *neural* interfaces which seek to deliver the full suite of stimuli directly to and from the brain [Charles 1999; Gosselin and Sawan 2010; Cozzi et al. 2005]. While this technology is extremely nascent and preliminary, it may represent the best potential combination of high-fidelity, high-detail delivery of sensory stimuli with low device intrusiveness.

3.1.6 The Uncanny Valley Revisited. Our discussion of realism in virtual worlds concludes with *the uncanny valley*, a term coined by Masahiro Mori to denote an apparent discontinuity between the additive likeness of a virtual character or robot to a human being and an actual human being’s reaction to that likeness. As seen

in Figure 1,³ the recognized humanity of a character or robot appears to increase steadily up to a certain point, at which human reactions turn negative or eerie at best, before “full” human likeness is attained [Mori 1970]. The phenomenon, observable in both computer-generated characters and humanoid physical devices, suggests that recognizing “full humanity” may involve a diverse collection of factors, including sensory subtleties, emotional or empathetic tendencies, diverse behaviors, and perceptual mechanisms for processing faces [Seyama and Nagayama 2007; Mieselhorn 2009; Seyama and Nagayama 2009; Tinwell et al. 2011].




Fig. 1. Masahiro Mori’s uncanny valley.

While the term is frequently seen in the context of video games, computer animation, and robotics, the uncanny valley takes on a particular significance for virtual world realism, since this form of realism involves psychological immersion, driven primarily by how we perceive and interact with avatars. As we become more proficient at rendering certain types of realism, we then have to try to understand which components of that realism have the greatest psychological impact, both individually and in concert with each other (the so-called multimodal approaches mentioned previously). Factors such as static facial detail vs. animation, as well as body shape and movement, have been studied for their ability to convey emotional content or to elicit trust [Hertzmann et al. 2009; McDonnell et al. 2009; McDonnell and Breidt 2010]. Some studies suggest that the uncanny valley response may also vary based on factors beyond the appearance or behavior of the human facsimile, such as by the gender, age, or even technological sophistication of the human observers [Ho et al. 2008; Tinwell and Grimshaw 2009].

As progress has been made in studying the phenomenon itself, so has progress also been made in crossing or transcending it. The computer graphics advances presented in this section have contributed to this progress in general, with research

³Based on a translation by Karl F. MacDorman and Takashi Minato, manually rendered and uploaded to Wikimedia Commons [Smurrayinchester (AV) and Voidvector (AV) 2008]

in facial modeling and rendering leading the way in particular. Proprietary algorithms developed by Kevin Walker and commercialized by Image Metrics use computer vision and image analysis of pre-recorded video from human performers to generate virtual characters that are recognizably lifelike and can credibly claim to cross the uncanny valley [Image Metrics 2011]. The company’s Faceware product has seen effective applications in a variety of computer games and animated films, as well as research that specifically targets the uncanny valley phenomenon [Alexander et al. 2009]. These techniques are still best applied to pre-rendered animations, with dynamic, real-time synthesis posing an on-going challenge [Joly 2010]. Alternative approaches include the modeling of real facial muscular anatomy [Marcos et al. 2010] and rendering of the entire human form, for applications such as sports analysis or telepresence [Vignais et al. 2010; Aspin and Roberts 2011].

For virtual world avatars to cross the uncanny valley, the challenge of effective visual rendering is compounded by (1) the need to achieve this rendering fully in real time, and (2) the use of real-time input (e.g., camera-captured facial expressions, synchronized voice chat, appropriate haptic stimuli). In these respects, virtual world environments pose the ultimate uncanny valley challenge, demanding *bidirectional*, *real-time* capture, processing, and rendering of recognizably lifelike avatars. Such avatars, in turn, are instrumental in producing the most immersive, psychologically compelling experience possible.

3.2 Ubiquity

The notion of *ubiquity* in virtual worlds derives directly from the prime criterion that a fully-realized Metaverse must provide a milieu for human culture and interaction that, as with the physical world, is psychologically compelling for the user. The real world is ubiquitous in a number of ways: first, it is *literally* ubiquitous — we unavoidably dwell in, move around, and interact with it, at all times, in all situations. Second, our *presence* within the real world is *ubiquitously manifest* — that is, our identity and persona are, under normal circumstances, universally recognizable, primarily through our physical embodiment (face, body, voice, fingerprints, retina) but augmented by a small, universally-applicable set of artifacts such as our signature, key documents (birth certificates, passports, licenses, etc.), and identifiers (social security numbers, bank accounts, credit cards, etc.). Our identity is further extended by what we *produce* and *consume*: books, music, or movies that we like; the food that we cook or eat; and memorabilia that we strongly associate with ourselves or our lives [Turkle 2007].

While our *perception* of the real world may sometimes fluctuate — such as when we are asleep — the real world’s literal ubiquity persists, regardless of what our senses say (and whether or not we are even alive). Similarly, while our identities may sometimes be doubted, inconclusive, forged, impersonated, or stolen, in the end these situations are borne of errors, deceptions, and falsehoods: there always remains a *real* “me,” with a core package of artifacts (physical, documentary, etc.) that authoritatively represent “me.”

To serve as a rich, alternative venue for human activity and interaction, virtual worlds must support some analog of these two aspects of real-world ubiquity. If a virtual world does not retain some constant presence and availability, then it feels remote and far-removed, and thus to some degree not as “real” as the physical

world. If there are barriers, artificial impedances, or undue inconveniences involved with identifying ourselves and the information that we create or use within or across virtual worlds, this distances our sense of self from these worlds and we lose a certain degree of investment or immersion within them. We thus divide the remainder of this section between these two aspects of ubiquity: ubiquitous availability and access, and ubiquitous persona and presence.

3.2.1 *Availability and Access of Virtual Worlds.* Ubiquitous availability of and access to virtual worlds can be viewed as a specific instance of *ubicomp* or *ubiquitous computing*, as first proposed by Mark Weiser [Weiser 1991; 1993]. The general definition of ubiquitous computing is the massive, pervasive presence of digital devices throughout our real-world surroundings. Ubiquitous computing applications include pervasive information capture, surveillance, interactive rooms, crowdsourcing, augmented reality, biometrics and health monitoring, and any other computer-supported activity that involves multiple, autonomous, and sometimes mobile loci of information input, processing, or output [Jeon et al. 2007].

The pervasive device aspect of ubicomp research — embedded systems, audio/video capture, ultramobile devices — can be of great service to the virtual worlds realm, especially in terms of a world’s ubiquitous availability. Physical access to virtual worlds has predominantly been the realm of traditional computing devices: PCs and laptops, ideally equipped with headsets, microphones, high-end graphics cards, and broadband connectivity. Visiting virtual worlds involves physically locating oneself at an appropriate station, whether a fixed desktop or a suitable setting for a laptop. Peripheral concerns such as connecting a headset, finding a network connection, or ensuring an appropriate aural atmosphere (i.e., no excessive noise) frequently arise, and it is only after one is “settled in” that entry into the virtual environment may satisfactorily take place.

While other features, such as realism and scalability (Section 3.4), may continue to be best served by traditional desktop and laptop stations, the current near-*exclusivity* of such devices as portals into virtual worlds hinders the perceived ubiquity of such worlds. When one is traveling, away from one’s “usual” system, or is in an otherwise “unwired” real-world situation, the inability to access or interact with a virtual world diminishes that world’s capacity to serve as an alternative milieu for cultural, social, and creative interaction — the very measure by which we differentiate virtual worlds from other 3D, networked, graphical, or social computer systems. Ubiquitous computing lays the technological roadmap that leads away from this current “desktop and laptop hegemony” toward pervasive, versatile availability and access to the Metaverse.

The latest wave of mobile and tablet devices comprises one aspect of the ubiquitous computing vision that has recently hit mainstream status. Combined with existing alternative access points to virtual worlds, such as e-mail and text-only chat protocols, these devices form the beginnings of ubiquitous virtual world access [Pocket Metaverse 2009]. Going forward, the increasing computational and audiovisual power of such devices may allow them to run viewers that approach current desktop/laptop viewers in fidelity and immersiveness. In addition, the new form factors and user interfaces enabled by such devices, with front-facing cameras, accelerometers, gyroscopes, and multitouch screens, should also facilitate new

forms of interaction with virtual worlds beyond the keyboard-mouse-headset combination that is almost universally used today [Baillie et al. 2005; Francone and Nigay 2011]. Such technologies make the ubiquitous availability of virtual worlds more compelling and immersive, in a manner that is different from traditional desktop/laptop interaction.

In parallel, modern web applications and web browsers, with their full programmability, standardized networking, 3D, and media interfaces, and zero-overhead installation or access mechanisms, comprise a software-side enabler toward virtual world ubiquity. Modern web architectures and browsers have already facilitated the ubiquity of e-commerce, social networking, and other applications. Virtual worlds can also attain this degree of availability as web standards, in conjunction with the platforms on which they run, begin to fulfill the requirements of a 3D virtual environment. The potential of such a platform has seen early demonstrations in systems like the Aves engine [Bakaus 2010b; 2010a]. Aves graphics were restricted to isometric 2D, but there is no reason to doubt that full 3D rendering using WebGL will eventually arrive.⁴

It should be noted that not all “browser-based” virtual environments are entirely browser-*native*, requiring third-party plugins to deliver an immersive experience. Such implementations only partially fulfill the ubiquity potential of a browser-based virtual world viewer, and may require significant porting or redesign to become browser-based in the true sense of the word.

In the end, the broad permeation of real-world portals into a virtual world, along the lines envisioned by ubiquitous computing, will elevate a virtual world’s ability to be an alternative and equally rich environment to the physical world: a true parallel existence to and from which we can shuttle freely, with a minimum of technological or situational hindrances. The current candidate platform for this vision of ubiquitous availability and access is a modern web application built on open-standard, browser-native technologies such as HTML5, Ajax, and WebGL, which, through its reliance solely on a web browser without third-party plug-ins, can then run on the full spectrum of devices with standards-compliant browsers.

3.2.2 *Manifest Persona and Presence.* As mentioned, in the real world we have a unified presence, centered on our physical body but otherwise permeating other locations and aspects of life through artifacts or credentials that represent us. Bank accounts, credit cards, memberships in various organizations, and other affiliations carry our manifest persona, the sum total of which contributes to our overall presence in society and the world.

This aggregation of data pertaining to our identity has begun to take place within personal, social, medical, financial, and other information spaces. The strength of this data aggregation varies, ranging from seamless integration (e.g., between two systems that use the same electronic credentials) to total separation within the digital space, only integrated outside via user input or interaction. If current trends continue, these informational fragments of identity should tend toward consolidation and interoperability, producing a distributed, ubiquitous electronic presence

⁴The Aves project was acquired by Zynga and has been brought completely in-house. However, the prospect of an open or freely available engine of this type remains likely and bright.

[Gilbert et al. 2011; Gilbert and Forney 2011] (in press). This form of ubiquity must also be modeled by virtual environments as a condition for their being a sufficient alternative existence to the real world.

As information systems have evolved from passive recorders of transactions, with the types of transactions themselves evolving from financial activity to interpersonal communication, documents and audiovisual media, and finally to social outlets, the flow of such information has also shifted from separated producers and consumers to so-called *prosumers* that simultaneously create, view, and modify content.⁵ A persona, in this context, is the virtual aggregate of a person’s online presence, in terms of the digital artifacts that he or she creates, views, or modifies.

The prosumer culture has entrenched itself well in blogs, mashups, and audiovisual and social media. Their close association to the individuals that produce and consume this content has provided greater motivation toward consolidation of persona than traditional information “wells” such as finance and commerce. Common information from social networks can be utilized across systems through published protocols and services, resulting in increasing unification and consolidation of one’s individuality [Facebook 2011; Twitter 2011; Yahoo! Inc. 2011; OAuth Community 2011; OpenID Foundation 2011]. This interconnectivity and transfer of persona, when and where it is required by the user, has yet to see adoption or implementation within virtual worlds. Virtual worlds are currently balkanized information silos, with the credentials required being just as disjoint and uncoordinated.

Shared credentials and the bidirectional flow of current information types such as writings, images, audio, and video across virtual worlds represent only the initial phases of bringing a ubiquitous, unified persona and presence into this domain. The actual *content* that can be produced/consumed within virtual environments expands as well, now covering 3D models encapsulating interactive and physics-based behaviors, as well as *machinima* presentations delivering self-contained episodes of virtual world activity [Morris et al. 2005]. Where interoperability (covered in the next section) concerns itself with the development of sufficiently extensible standards for conveying or transferring such digital assets, ubiquity concerns itself with consistently *identifying* and *associating* such assets with the personas that produce or consume them. Virtual worlds must provide this aggregation of virtual artifacts under a unified, omnipresent electronic “self” in order to serve as a viable real-world alternative existence. Digital assets which comprise a virtual self must also be seamlessly available from all points of virtual access, as illustrated in Figure 2.

To date, no system facilitates ease of migration for the virtual environment persona across virtual worlds, unlike the interchange protocols in place for social media outlets and a host of other information hosting and sharing services. There are some efforts to create unified standards for virtual environments [IEEE VW Standard Working Group 2011a; COLLADA Community 2011; Arnaud and Barnes 2006; Smart et al. 2007]. At present, these efforts have yet to bear long-term fruit.

In order to truly offer a ubiquitous persona in virtual environments, the individual must be able to seamlessly and transparently cross virtual borders while remaining connected to existing, legally-accessible information assets. A solution

⁵The term “prosumer” here combines the terms “producer” and “consumer,” and not “professional consumer,” as the term is used in other domains.

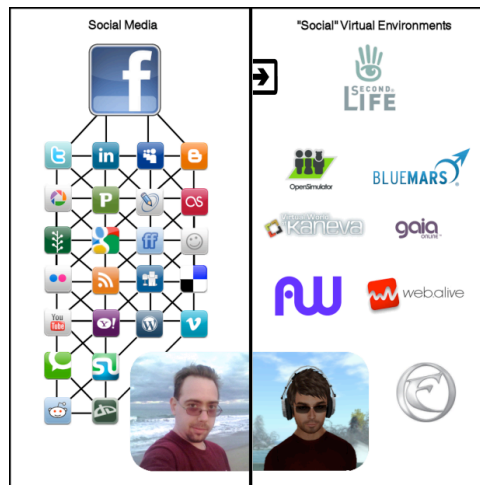


Fig. 2. Interactions (or lack of) among social media networks and virtual environments.

for this requirement may involve not only industry consensus but also some technical innovation that can connect these systems and store assets securely for transfer while remaining independent of any single central influence.

3.3 Interoperability

In terms of functionality alone, interoperability in the virtual worlds context is little different from the general notion of interoperability: it is the ability of distinct systems or platforms to exchange information or interact with each other seamlessly and, when possible, transparently. Interoperability also implies some type of consensus or convention, which, when formalized, then become standards.

As specifically applied to virtual worlds, interoperability might be viewed as merely the enabling technology required for ubiquity, as described in the previous section. While this is not inaccurate, interoperability remains a key feature of virtual worlds in its own right, since it is interoperability that puts the capital *M* in the Metaverse: just as the singular, capitalized Internet is borne of layered standards which allow disparate, heterogeneous networks and subnetworks to communicate with each other transparently, so too will a singular, capitalized *Metaverse* only emerge if corresponding standards also allow disparate, heterogeneous *virtual worlds* to seamlessly exchange or transport objects, behaviors, and avatars.

This desired behavior is analogous to travel and transport in the real world. As our bodies move between physical locations, our identity seamlessly transfers from point to point, with no interruption of experience. Our possessions can be sent from place to place and, under normal circumstances, they do not substantially change in doing so. Thus, real-world travel has a *continuity* in which we and our things remain largely intact in transit. We take this continuity for granted in the real world, where it is indeed a matter of course. With virtual worlds, however, this effect ranges from highly disruptive to completely non-existent.

The importance of having a single Metaverse is connected directly to the long-

term endpoint of having virtual worlds offer a milieu for human sociocultural interaction that, like the physical world, is psychologically rich and compelling. Such integration immediately makes all compatible virtual world implementations, regardless of lineage, parts of a greater whole. With interoperability, especially in terms of a transferable avatar, users can finally have full access to any environment, without the disruption of changing login credentials or losing one’s chain of cross-cutting digital assets.

3.3.1 Existing Standards. The first [relatively] widely-adopted standard relating to virtual worlds is VRML (Virtual Reality Modeling Language). Originally envisioned as a 3D-scene analog to the then-recently-emergent HTML document standard, VRML captured an individual 3D space as a single “world” file that can be downloaded and displayed with any VRML-capable browser. VRML, like HTML, utilized URLs to facilitate “navigation” from one “world” to another, thus enabling an early form of virtual world interoperability.

VRML has since been succeeded by X3D. X3D expands on the graphics capabilities of VRML while also adjusting its syntax to align it with XML. Both formats reached formal, ISO recognition as standards, but neither has gained truly mainstream adoption. Barriers to this adoption include the lack of a satisfactory, widely-available client for effectively displaying this content, competition from proprietary efforts, and, most importantly, a lack of critical mass with regard to actual virtual world content and richness. Efforts are currently under way for making X3D a first-class citizen of the HTML5 specification, in the same way that other XML dialects like SVG and MathML have been integrated. Such inclusion may solve the issue of client availability, as web browser developers would have to support X3D if they are to claim full HTML5 compliance, but overcoming the other barriers — especially that of insufficient content — would require additional effort and perhaps some degree of advocacy.

Technologically distinct and more recent than VRML and X3D is COLLADA [COLLADA Community 2011; Arnaud and Barnes 2006]. COLLADA has a different design goal: it is not as much a complete virtual world standard as it is an *interchange* format. Widespread COLLADA adoption would facilitate easier exchange of objects and behaviors from one virtual world system to another; in fact, COLLADA’s scope is not limited to virtual worlds, as it can be used as a general-purpose 3D object interchange mechanism.

On the opposite side of the standards aisle is the *de facto* “standard” established by Linden Lab, which stems purely from the relative success of Second Life. Unlike X3D and VRML before it, Second Life crossed an adoption threshold that spawned a fairly viable “ecology” of alternative viewers and servers. Beyond this widespread use, however, Second Life user credentials, protocols, and data formats fall short. They remain proprietary, with an array of restrictions on how and where they can be used. For instance, only viewers (clients) have official support from Linden Lab; alternative *server* implementations such as OpenSimulator are actually reverse-engineered and not directly sanctioned.

As with other areas with interoperability concerns, virtual world interoperability is not an entirely technical challenge, as many of its current players actually benefit from a lack of interoperability in the short term and thus understandably resist

it or are apathetic to it. Similar to other areas (such as the emergence of the Worldwide Web above proprietary online information services), a potential push toward interoperability requires multilateral contributions and momentum from standards groups, commercial developers, and the open source community.

3.3.2 Virtual World Interoperability Layers. The promise of virtual world interoperability as envisioned in this paper — the kind that facilitates a single Metaverse and not a disjointed, partitioned scattering of incompatible environments — is currently a work in progress. Ideally, such a standard would have the rigor and ratification achieved by VRML and X3D, while also attaining a critical mass of active users and content. In addition, this standard should actually be a *family* of standards, each concerned with a different layer of virtual world technology:

- (1) A *model* standard should sufficiently capture the properties, geometry, assets, and perhaps behavior of virtual world environments. VRML, X3D, and COLLADA are predominantly model standards. Note that a model standard’s main purpose is to facilitate object interchange across different virtual world systems. The systems themselves may represent objects using a private or proprietary mechanism, as long as they seamlessly read from and write to the non-proprietary, open standard. Modeling is the most obvious area that needs standardization, but this alone will not facilitate the capital-*M* Metaverse.
- (2) A *protocol* standard defines the interactive and transactional contract between a virtual world client or viewer and a virtual world server. This layer enables a mix-and-match situation with multiple choices for virtual world clients and servers, none of which would be “wrong” from the perspective of compatibility. As mentioned, Second Life has established a de facto protocol standard solely through its widespread use. Full standardization, at this writing, does not appear to be part of this protocol’s road map.
Open Cobalt currently represents the most mature effort, other than Second Life, that can potentially fulfill the role of a virtual world protocol standard [Open Cobalt Project 2011]. As an open source platform under the flexible MIT free software license, Open Cobalt certainly offers fewer limits or restrictions to widespread adoption. Like VRML and X3D, however, the protocol will need connections to diverse and compelling content in order to gain critical mass.
- (3) A *locator* standard for identifying places or landmarks across virtual worlds. This may be, technologically, the easiest standard to establish, since it already exists: uniform resource locators (URLs), a subset of the uniform resource identifiers (URI) standard [Berners-Lee et al. 2005], are completely adaptable for virtual world landmarks, and have in fact been used by Linden Lab for Second Life locations (SLURLs, or Second Life URLs). The missing piece is universal consensus and adoption of a virtual world URL scheme (VWURLs, so to speak, or MVURLs if the preferred phrase is “Metaverse URL”) so that, regardless of the virtual world system or viewer, the same scheme can be recognized and used to move easily from virtual place to virtual place.
- (4) An *identity* standard defines a unified set of user credentials and capabilities that can cross virtual world platform boundaries: this would be the virtual world equivalent of efforts such as OpenID and the Facebook Connect [OpenID

Foundation 2011; Facebook 2011]. The importance of an identity standard was covered in Section 3.2.2.

- (5) A *currency* standard quantifies the value of virtual objects and creations, facilitating their trade and exchange. Linden Dollars within Second Life represent the dominant virtual world currency, but it can be used only within Second Life and is meaningful only within that system (until exchanged for a real-world currency). Bitcoin is not controlled by any particular authority, with a number of mechanisms in place for tracking or validating its transactions while retaining a peer-to-peer nature [Nakamoto 2009; Piotrowski and Terzis 2011]. A standard such as this may form the basis for interoperable Metaverse currency.

As with Internet standards, the layering is defined only to separate the distinct virtual world concerns, to help focus on each area’s particular challenges and requirements. Ultimately, standards at *every* layer will need to emerge, with the combined set finally forming an overall framework around which virtual worlds can connect and coalesce into a single, unified Metaverse, capable of growing and evolving in a parallel, decentralized, and organic manner.

3.4 Scalability

As with the features discussed in the previous sections, virtual worlds have scalability concerns that are similar to those that exist with other systems and technologies while also having some distinct and unique issues from the virtual world perspective. Based on this paper’s prime criterion that a fully-realized Metaverse must provide a milieu for human culture and interaction, scalability may thus be the most challenging virtual world feature of all, as the physical world is of *enormous* and potentially *infinite* scale on many levels and dimensions. Three dimensions of virtual world scalability have been identified in the literature [Liu et al. 2010]:

- (1) *Concurrent users/avatars* — The number of users interacting with each other at a given moment
- (2) *Scene complexity* — The number of objects in a particular locality, and their level of detail or complexity in terms of both behavior and appearance
- (3) *User/avatar interaction* — The type, scope, and range of interactions that are possible among concurrent users (e.g., intimate conversations within a small space vs. large-area, crowd-scale activities such as “the wave”)

In many respects, the virtual world scalability problem parallels the computer graphics rendering problem: what we see in the real world is the constantly updating result of multitudes of interactions among photons and materials governed by the laws of physics — something that, ultimately, computers can only *approximate* and never absolutely *replicate*. Virtual worlds add further dimensions to these interactions, with the human social factor playing a key role (as expressed in the first and third dimensions above). Thus it is no surprise that most of the scientific problems in [Zhao 2011] focus on theoretical limits to virtual world modeling and computation, since, in the end, what else is a virtual world but an attempt to simulate the real world *in its entirety*, from its physical makeup all the way to the activities of its inhabitants?

3.4.1 *Traditional Centralized Architectures.* Virtual world architectures, historically, can be viewed as the progeny of graphics-intensive 3D games and simulations, and thus they share the strengths and weaknesses of both [Waldo 2008]. Graphics-intensive 3D games focus on scene complexity and detail, particularly in terms of object rendering or appearance, and strongly emphasize real-time performance. Simulations focus on interactions among numerous elements (or *actors*), with a focus on repeatability, accuracy, and precision.

Both 3D games and simulations, along with many other application categories at the time, were predominantly single-threaded and centralized when the first virtual world systems emerged. Processing followed a single, serialized computational path, with some central entity serving as the locus of such computations, the authoritative repository of system state, or both. Virtual world implementations were no exception; as 3D-game-and-simulation hybrids, however, they consisted of *two* “centers:” clients or *viewers* that rendered the virtual world for its users and received the actions of its avatars (derived from 3D games), and servers that tracked the shared virtual world state, changing it as avatar and scripted directives were received from the clients connected to them (derived from simulations).

This initial paradigm was in many respects unavoidable and quite natural, as it was the reigning approach for most applications at the time, including the 3D game and simulation technologies that converged to form the first virtual world systems. However, this initial paradigm also established a form of technological inertia, with succeeding virtual world iterations remaining conceptually centralized despite early evidence that such an architecture ultimately does not scale along any one of the aforementioned dimensions (concurrent avatars, scene complexity, avatar interactions), much less all three [Lanier 2010; Morningstar and Farmer 1991]. This “technology lock-in” resulted in successive attempts at virtual world scalability that addressed the *symptoms* of the problem, and not its ultimate cause [Lanier 2010].

3.4.2 *Initial Distribution Strategies: Regions and Shards (Distribution by Geography).* The role of clients/viewers as renderers of an individual user’s view of the virtual world was perhaps the most obvious initial boundary across which computational load can be clearly distributed: the same geometric, visual, and environmental information can be used by multiple clients to compute differing views of the same world in parallel. Outside of this separation, the maintenance and coordination of this geometric, visual, and environmental data — the *shared state* of the virtual world — then became the initial focus of scalability work.

Initial approaches to distributing the maintenance of shared state fall within two broad categories: *regions* and *shards*. Both distribution strategies are based on the virtual world’s “geography.” In the *region* approach, the total expanse of a virtual world is divided into contiguous, frequently uniform chunks, with each chunk being given to a distinct computational unit or *sim*. This division is motivated by an assumption that activities which affect each other are most likely to take place within a relatively small area (i.e., *spatial locality*). Spatial locality implies that an individual avatar is primarily concerned with or affected by activities within a finite radius of that avatar. The “scaling” rationale thus sees the growth of a virtual world as the accretion of new regions, reflected by an increase in active sims. Individual avatars are then serviced by only one computational unit or host

at a time because events of concern are thought to be restricted to within a small radius of that avatar.

The *shard* approach is also based on a virtual world’s “geography,” with portions of that geography being *replicated* rather than partitioned across multiple hosts or servers. Shards prioritize concurrent users over a single world state. Since parts of the virtual world are effectively copied across multiple servers, more users can seem to be located in a given area at a given time. Tradeoffs for this concurrency include not being able to interact with *every* user in the same area (since users are separated over multiple servers). Most severely, as mentioned, the virtual world loses its single shared state. Thus, shards are applicable primarily to transient virtual environments such as massively multiplayer online games, and not to general virtual worlds that must maintain a single, persistent, and consistent state.

The region approach has been the de facto distribution strategy for many virtual world systems, including Second Life [Wilkes 2008]. However, as pointed out by Ian Wilkes and borne out by subsequent empirical studies [Gupta et al. 2009; Liu et al. 2010], partitioning by region fails on key virtual world use cases:

- (1) Number of concurrent users in a single region — even the most powerful sims saturate at around 100 concurrent users (and thus a potential of 4950 concurrent user-to-user interactions, not to mention interactions with objects)
- (2) Varying complexity and detail in regions — the spatial basis for regions does not take into account the possibility that one region may be relatively barren and inactive, while another may include tens of thousands of objects with multiple concurrent users
- (3) Scope of user interaction — virtual world users do not necessarily interact solely with users within the same region; services such as instant messaging and currency exchange may cross region boundaries
- (4) Activity at region boundaries — since region boundaries are actually transitions from one sim (host) to another, activities across such boundaries (travel, avatar interaction) incur additional latency due to data/state transfer

In the case of Second Life [Wilkes 2008], the third use case is addressed through the creation of separate services distinct from the grid of region sims. Such services include user/login management, voice chat, and raw data or asset storage. The separation of these subsystems did lighten the load on regions somewhat, but only for very specific types of load.

The fourth case is addressed by policy, through avoiding adjacent regions when possible. The first and second cases remain unresolved beyond the use of increasingly powerful individual servers and additional network bandwidth. OpenSimulator, as a reverse-engineered port of the Second Life architecture, thus also manifests the same strengths and weaknesses of that architecture [OpenSimulator Project 2011]. It has, however, also served as a basis for more advanced work on virtual world scalability [Liu et al. 2010].

3.4.3 Distribution by Regions and Users. Future iterations of Second Life propose to separate computation and information spaces into two distinct dimensions: separation of user/avatar state from region state, and the establishment of multiple such user and region grids [Wilkes 2008]. The overall interaction space and

computational load are thus divided by two linear factors: user/avatar vs. region interactions, and distinct clusters of users/avatars and regions. It is currently unknown whether this transition has been accomplished by Linden Lab.

This direction, particularly the accommodation of multiple grids of user and region servers and the ability to add, remove, or reallocate distinct computational units for handling user and world interactions, resembles a peer-to-peer (P2P) architecture in topology. However, unless interfaces for these units are cleanly abstracted and published, and until new instances of user or region servers are permitted to start/stop on any Internet-connected host, this projected approach remains strictly within the boundaries of Linden Lab.

The Solipsis and OpenCobalt projects do not have such boundaries, however, and so can be viewed as P2P architectures in terms of both technology and policy [Keller and Simon 2002; Frey et al. 2008; Open Cobalt Project 2011]. Both systems allow instances to dynamically augment the virtual world network’s geography (through *nodes* in Solipsis and *spaces* in OpenCobalt) and user space (a user’s running peer serves as the origin for that user’s avatar within the virtual world). Solipsis uses neighborhood algorithms to locate and interact with other nodes, while OpenCobalt employs an LDAP registry to “publish” available spaces, with private spaces remaining accessible if a user knows typical network parameters such as an IP address and port number. Interactions in Solipsis are scoped through an area of interest (AOI) approach (*awareness radius*) while OpenCobalt employs reliable, replicated computation based on David Reed’s TeaTime protocol [Open Cobalt Project 2011; Reed 2005].

3.4.4 Distribution by Region State, Users, and State-Changing Functions. With the widely acknowledged limitations of existing virtual world architectures, on-going work on virtual world scalability centers on additional axes of computational distribution, as well as the overall system topology of multiple virtual worlds (“metagalaxies”). At this point, no single approach has yet presented, publicly, both a stable, widely-used reference implementation alongside empirical, quantitative evidence of significant scalability improvement. A fundamental issue behind any approach to virtual world scalability is that concurrent interactions are unavoidably quadratic: for any n concurrently interacting avatars, agents, or objects — whether concurrent by location, communication, or any other relationship — there can always be on the order of n^2 updates among that cluster of participants. Approaches to addressing this quadratic growth range from *constraints* on the possible interactions or number of interacting objects, to *replicating* computations so that the same result can be reached by multiple computational units with minimal messaging.

The RedDwarf Server project, formerly Project Darkstar, takes a converse approach to scalability by allowing clients a *perceived* single-threaded, centralized system, where the “server” is itself a container for distributed synchronization, data storage, and dynamic reallocation of computational tasks [Waldo 2008; RedDwarf Server Project 2011]. The project itself is not exclusively targeted for virtual worlds, with online games and non-immersive social networks as possible applications as well. Returning briefly to online games, Trion Worlds’s *Rift* is noteworthy here because its server architecture also claims to use a similar, dynamic reallocation approach as Project Darkstar, along with the distribution of computational

load by functions beyond shards [Gladstone 2011; Gamasutra Staff 2011].

SEVE (Scalable Engine for Virtual Environments) uses an *action-based protocol* that leverages the *semantics* of virtual world activities to both streamline the communication between servers and clients as well as optimistically replicate state changes across all active hosts [Gupta et al. 2009]. Virtual world agents generate a queue of *action–result* pairs, with actions sent to the server as they are performed at each client. The role of the server is to reconcile the actions of multiple agents into a single, authoritative “state of the world.” Successive changes to the state of the world are redistributed to affected clients periodically, with adjustments made to each client’s rendering of the world as needed. SEVE has been shown to scale well in terms of concurrent clients and complex actions, though individual clients themselves may require additional power and storage because they replicate the world state while running.

The *Distributed Scene Graph* (DSG) approach separates the state of the virtual world into multiple *scenes*, with computational *actors* each focused on specific types of changes to a scene (e.g., physics, scripts, collision detection) and communication-intensive *client managers* serving as “liaisons” between connected clients/viewers and the virtual world [Liu et al. 2010]. This separation strategy isolates distinct computational concerns from the data/state of the world and communications among clients that are “subscribed” to that world. This approach appears to be very amenable to dynamic load balancing, since specific types of work (e.g., physics calculations, script execution, avatar interaction) can be distributed among multiple machines, with the “state of the scene” serving as the unifying target of each server’s work. Liu et al. are also experimenting with dynamic level-of-detail reduction, allowing viewers that “see” broad areas of a virtual world to receive and process less scene data — an optimization suggested by work on realism and perception [Blascovich and Bailenson 2011; Glencross et al. 2006]. The system is implemented on top of OpenSim.

3.4.5 Proposed Architectural Direction. The constraints and requirements of virtual world scalability, including the need to maximize distribution in the face of limited bandwidth, have constituted a long-standing challenge to the field [Morningstar and Farmer 1991]. All told, while no single system has yet established itself as a clear next step, the components needed for such a system appear to have emerged among the various works surveyed in this section.

In terms of an overall strategy, the Distributed Scene Graph (DSG) architecture offers the broadest degree of distribution, the performance of which has been supported by an OpenSim-based prototype and preliminary studies [Liu et al. 2010]. The core of this architecture, which is a set of scene servers whose sole responsibility is to manage the state of a specific region within the virtual world, may best be implemented using a system such as RedDwarf, whose features and strengths correspond well with the role of a DSG scene server. With RedDwarf-driven scene servers or grids, scene management itself can scale internally, based on the size and activity within a particular section of the virtual world. While no specific recommendations can be made for the other components of the DSG strategy at this time — physics engines, scripting machines, user managers, etc. — the very separation of other virtual world functions across these broad categories facilitates

open innovation and experimentation in these areas.

Assuming that a standard “scene transaction” protocol exists for scene servers (whether DSG’s, RedDwarf’s, or any other functionality sufficient specification), another key architectural element would be the dynamic, open-ended capabilities afforded by the peer-to-peer (P2P) approach taken by Solipsis and OpenCobalt. P2P complements DSG very well in a variety of ways: first, P2P facilitates open growth of a virtual world system by allowing new scene servers (and thus new places) to join it at any time, from anywhere on the Internet. An architecture that includes the replicated computation functionality of OpenCobalt will also allow *individual* scenes to accrue additional servers as needed, for particularly busy or heavily-populated periods.

DSG’s distribution across multiple axes (region, user, function) also benefits greatly from P2P. With the separation of scripting servers from scene servers, for example, P2P will permit users’ own machines, or others, to participate in the execution of scripts, thus distributing that computational load quite effectively. The ability to add more client managers as a particular concurrent user population grows addresses that particular axis of virtual world scalability. As mentioned earlier, P2P issues are not entirely technical; successful establishment of P2P in virtual worlds will require that no artificially restrictive policies prevent other systems from joining or contributing to the Metaverse as needed.

Where the DSG strategy addresses the distribution of computational load, the *communication* load, which increases quadratically as the number of peers increases, can be addressed by area of interest (AOI) or other clustering techniques, or by reducing or coalescing the information that is exchanged by servers, as seen in SEVE’s semantic, action-based protocol and others. This final piece completes the overall proposed blueprint for a fully scalable Metaverse architecture, illustrated in Figure 3 and derived of Figure 5 from [Liu et al. 2010].

4. SUMMARY AND CONCLUSIONS

In the past three decades considerable progress has been made in moving from text-based multi-user virtual environments to the technical implementation of advanced virtual worlds that previously existed only in the literary imagination. Contemporary virtual worlds are now complex immersive environments with increasingly realistic 3D graphics, integrated spatial voice, content creation tools, and advanced economies. These progressive capabilities enable them to serve as elaborate contexts for work, socialization, creativity, and play and to increasingly operate more like digital cultures than as games.

Virtual world development now faces a major new challenge: how to move from a set of sophisticated, but completely independent, immersive environments to a massive, integrated network of 3D virtual worlds or Metaverse and establish a parallel context for human interaction and culture. The current work has defined success in this effort as achieving significant progress with regard to four features that are considered central elements of a fully-realized Metaverse: realism (enabling users to feel fully immersed in an alternative realm), ubiquity (establishing access to the system via all existing digital devices and maintaining the user’s virtual identity throughout all transitions within the system), interoperability (allowing 3D objects

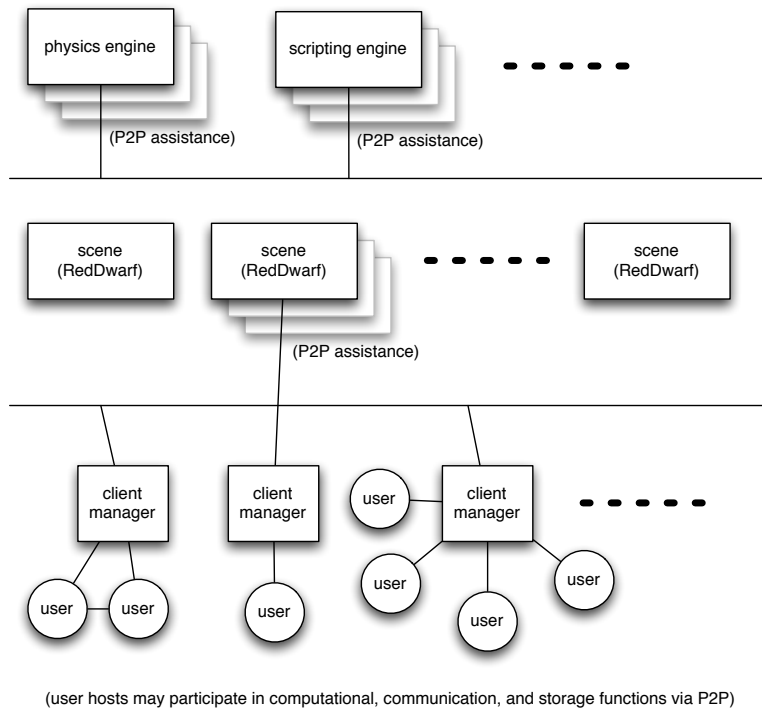


Fig. 3. Proposed architectural direction for a scalable Metaverse.

to be created and moved anywhere and users to have seamless, uninterrupted movement throughout the system) and scalability (permitting concurrent, efficient use of the system by massive numbers of users). The current status and needed developments to move to a fully functioning Metaverse with respect to each of these essential features can be summarized as follows:

Realism. Technologies and techniques for improving immersive realism, especially for the senses of sight, hearing, and touch, continue to evolve independently of virtual worlds due to their applications in other domains, such as games, entertainment, and social media. Bringing these developments into the Metaverse is typically a phase behind innovations in these other contexts due to the additional requirement for bidirectional information flow (i.e., input and output) and real-time performance. Moving forward, realism-related work must achieve either concurrent input/output capability or real-time performance, or both, in order to achieve a useful role in virtual world applications. In addition, developments for all senses must be integrated to form a cohesive, immersive whole, exemplified by the challenge presented by the uncanny valley. “Crossing” the uncanny valley serves as an endpoint toward which innovations in virtual world realism continues to strive and possibly lead, because of the greater impact that psychological engagement has on virtual world realism than, say, realism in visual effects or computer games.

Ubiquity. The first aspect of ubiquity — *ubiquitous availability and access* to virtual worlds — rides on the crest of developments in ubiquitous computing in general. As ubicomp has progressed, access to virtual worlds has begun to move beyond a stationary “desktop PC rig,” expanding now into laptops, tablets, mobile devices, and augmented reality. In addition, the evolution of web browsers from static document viewers into programmable 3D engines in their own right can potentially reduce the barrier to entry into a virtual world environment to a single, addressable link, without requiring time-consuming software installation or updates/maintenance. The second aspect — *ubiquitous identity*, or *manifest persona* — has emerged as multiple avenues of digital expression (blogs, social networks, photo/video hosting, etc.) have become increasingly widespread. These artifacts extrapolate naturally into virtual world artifacts such as virtual clothing, objects, and buildings but currently suffer from balkanization across multiple user accounts and information hubs that integrate only incidentally, through the individual or individuals who know the usernames and passwords to each of these venues.

Progress in ubiquitous availability and access to virtual worlds may very well take care of itself, as the increasing power, diversity, and availability of general computing devices, as well as the capabilities of browser-based implementations, continue to advance. As long as virtual world developments move alongside general ubicomp developments, moving in and out of the Metaverse may become as convenient and fluid as browsing the Worldwide Web is today. In contrast, achieving ubiquitous identity will require more concerted and coordinated effort, both technologically and socially/politically, to allow one’s information to flow more freely among various milieus. Social networks are beginning to move in this direction, though have not fully reached it yet, and virtual worlds must follow suit in this area. It is also possible that virtual worlds may take a leadership position in this regard, as virtual world artifacts may be more closely linked to one’s digital persona due to the immersive environment (and therefore may motivate progress in this area more than social networks or blog/community sites can).

Interoperability. Interoperability in virtual worlds currently exists as a loosely connected collection of information, format, and data standards, most of which focus on the transfer of 3D models/objects across virtual world environments. While many standards have been proposed (e.g., VRML, X3D, COLLADA), no single one has gained long-term traction (though COLLADA is showing some promise). However, virtual world interoperability is not solely limited to 3D object transfer: true interoperability also involves communication protocol, locator, identity, and currency standards (with the first two having direct equivalents in the Worldwide Web). The widespread adoption of standards, in any area of computing, is frequently borne of much work, collaboration, and advocacy, where the most difficult barriers may not actually be technological or computational but social or political. Accomplishing virtual world interoperability is no different, and in some ways is exacerbated by the sheer breadth of elements and layers that must transfer seamlessly across different virtual world systems. Candidate standards exist; however, they are either in primary use outside of virtual worlds, or else remain preliminary and are still very much works-in-progress. The wide-ranging requirements and scope of digital assets involved in virtual worlds have the potential of making the Metaverse

the “killer app” that finally leads the charge toward seamless interoperability.

Scalability. Virtual world technologies are currently in an initial stage of departure from highly centralized system architectures, with the distribution of load being predominantly focused on contiguous regions or areas. The limitations of this strategy have been identified and multiple independent initiatives have proposed potential next steps ranging from peer-to-peer approaches to divisions of computational and communication load in ways that address the particular scalability needs of virtual world environments [Liu et al. 2010]. Going forward, an integrative phase is needed where the multiple independent research threads are brought together in complementary and cohesive ways to form a whole that is greater than the sum of its parts. From the surveyed work, an overall architectural strategy based on DSG (with a high-performance “world state” system like RedDwarf at its core), peer-to-peer openness as seen in Open Cobalt and Solipsis for broad distribution of load (not an exclusively technological challenge), and reduction/optimization techniques like action-based protocols appears to hold the most promise for massive scaling of an integrated virtual world system. As echoed in the NAE Grand Challenges, this is scaling of the highest order: progress in virtual world scalability implies progress in the scalability of many other types of multiuser, multitiered systems.

There are several factors that promote optimism that a fully developed Metaverse can be achieved, as well as a number of significant constraints to realizing this goal. On the encouraging side, the National Academy of Engineering’s designation of enhanced virtual reality (with virtual worlds a major subset of this field) as one of the 14 grand challenges for engineering in the 21st century [National Academy of Engineering 2008], along with government reports from China and Japan identifying virtual reality as a priority technology for development [Zhao 2011] provides a strong impetus from major institutional sources for continued innovation in the field. Another source of encouragement comes from the continued expansion of public interest in the use of virtual worlds. According to Kzero, a virtual world monitoring service, the number of worldwide registered accounts in virtual worlds approached 1.4 billion as of the end of the second quarter of 2011, representing an overall growth rate of 338% from approximately 414 million global accounts in slightly over 2 years (<http://www.slideshare.net/nicmitham/kzero-universe-q2-2011>). Moreover, approximately 70% of current account holders are below the age of 16, with almost half (652 million) in the 10–15 year old segment, and only 3% of worldwide users over the age of 25. These data indicate that a new generation accustomed to graphically rich, 3D digital environments (both virtual worlds and immersive games offered online and through consoles such as PlayStation, Xbox360, and Wii) is rapidly coming of age and will likely fuel continued development in all immersive digital platforms including advanced virtual worlds.

Finally, continued improvement in hardware performance, as reflected in Moore’s law and Kurzweil’s law of accelerating returns [Moore 1965; Kurzweil 2001], promotes confidence that sufficient computing efficiency will be available to achieve real time performance of any computable function related to the operation of virtual

worlds. In sum, top-down forces (i.e., the priority focus of scientific and governmental institutions), bottom-up pressures stemming from expanding popular interest in immersive digital environments, and continued advances in hardware capacity and performance, all contribute to positive expectations for progress in achieving a fully-realized Metaverse.

Along with forces that are propelling the development of the Metaverse forward, there are two significant barriers that may inhibit the pace or extent of this progress. The first pertains to current limits in computational methods related to virtual worlds. It is noteworthy that most of the items contained in Zhao’s delineation of the major challenges facing virtual world development (such as determining whether it is possible to digitally model any given object in the physical world — a “theory of modelability” — or whether a duality/tradeoff truly exists between virtual environment fidelity and real-time performance) focus on computational issues (i.e., the development of algorithms, theories, concepts, and paradigms) rather than hardware problems [Zhao 2011]. In fact, based upon current understanding, the computability of many of these key processes is presently unknown and, thus, conceptual advances will play a pivotal role in future progress toward creating a viable Metaverse. For example, because most of the major algorithms related to immersive realism have been developed and established, significant progress in real-time performance and bidirectional information flow of elements that contribute to this feature is likely in the immediate future. In contrast, in the area of scaling for example, conceptual solutions related to the optimum architectural configuration are still being considered and, until such time as a consensus is achieved on this issue, it is unlikely that there will be rapid progress in achieving a massively scaled Metaverse.

In addition to conceptual and computational challenges, the development of the Metaverse may be constrained by significant economic and political barriers. Currently virtual worlds are dominated by proprietary platforms such as Second Life, Cryworld, Utherverse, IMVU, and World of Warcraft, or government-controlled worlds such as the China-based Hipihi. These platforms have played an important role in the history of virtual worlds by providing development capital and enhancing public awareness of the technology. However, just as the old walled gardens of AOL, CompuServe, and Prodigy were instrumental in expanding Internet usage early on, but ultimately became an inhibitory force in the development of the Worldwide Web, these proprietary and state-based virtual world platforms have sparked initial growth but now risk constraining innovation and advancement. As we have seen with the development of the Internet, progress has been best served by the combined participation and innovation of proprietary and open source initiatives. Similarly, the advancement of a fully-realized Metaverse would likely be maximized by harnessing the same process of collective effort and mass innovation that was instrumental in the creation and expansion of the Web.

REFERENCES

- ABATE, A. F., GUIDA, M., LEONCINI, P., NAPPI, M., AND RICCIARDI, S. 2009. A haptic-based approach to virtual training for aerospace industry. *J. Vis. Lang. Comput.* 20, 318–325.
- ACTION = REACTION LABS, L. 2010. Ghost dynamic binaural audio. <http://www.actionreactionlabs.com/ghost.php>. Accessed on March 7, 2011.

- ADABALA, N., MAGNENAT-THALMANN, N., AND FEI, G. 2003. Real-time rendering of woven clothes. In *Proceedings of the ACM symposium on Virtual reality software and technology*. VRST '03. ACM, New York, NY, USA, 41–47.
- ALEXANDER, O., ROGERS, M., LAMBETH, W., CHIANG, M., AND DEBEVEC, P. 2009. The digital emily project: photoreal facial modeling and animation. In *ACM SIGGRAPH 2009 Courses*. SIGGRAPH '09. ACM, New York, NY, USA, 12:1–12:15.
- ANTONELLI (AV), J. AND MAXSTED (AV), M. 2011. Kokua/Imprudence Viewer web site. <http://www.kokuaviewer.org>. Accessed on March 30, 2011.
- ARANGARASAN, R. AND PHILLIPS JR., G. N. 2002. Modular approach of multimodal integration in a virtual environment. In *Proceedings of the 4th IEEE International Conference on Multimodal Interfaces*. ICMI '02. IEEE Computer Society, Washington, DC, USA, 331–336.
- ARNAUD, R. AND BARNES, M. C. 2006. *COLLADA: Sailing the Gulf of 3D Digital Content Creation*. A K Peters Ltd, Natick, MA, USA.
- ASPIN, R. A. AND ROBERTS, D. J. 2011. A GPU based, projective multi-texturing approach to reconstructing the 3D human form for application in tele-presence. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. CSCW '11. ACM, New York, NY, USA, 105–112.
- AU, W. 2008. *The making of Second Life*. Collins, New York.
- AURORA-SIM PROJECT. 2011. Aurora-Sim web site. <http://www.aurora-sim.org>. Accessed on March 30, 2011.
- BAILLIE, L., KUNCZIER, H., AND ANEGG, H. 2005. Rolling, rotating and imagining in a virtual mobile world. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*. MobileHCI '05. ACM, New York, NY, USA, 283–286.
- BAKAUS, P. 2010a. Aves engine sneak preview. <http://www.youtube.com/watch?v=013qQ4CEUTo>. Accessed on July 12, 2011.
- BAKAUS, P. 2010b. Building a game engine with jQuery. jQuery Conference 2010, <http://www.slideshare.net/pbakaus/building-a-game-engine-with-jquery>. Accessed on July 12, 2011.
- BERNERS-LEE, T., FIELDING, R., AND MASINTER, L. 2005. Uniform resource identifier (URI): Generic syntax. internet STD 66, RFC 3986. <http://www.ietf.org/rfc/rfc2396.txt>. Accessed on July 12, 2011.
- BIOCCA, F., KIM, J., AND CHOI, Y. 2001. Visual touch in virtual environments: An exploratory study of presence, multimodal interfaces, and cross-modal sensory illusions. *Presence: Teleoper. Virtual Environ.* 10, 247–265.
- BLANCHARD, C., BURGESS, S., HARVILL, Y., LANIER, J., LASKO, A., OBERMAN, M., AND TEITEL, M. 1990. Reality built for two: a virtual reality tool. *SIGGRAPH Comput. Graph.* 24, 35–36.
- BLASCOVICH, J. AND BAILENSON, J. 2011. *Infinite Reality: Avatars, Eternal Life, New Worlds, and the Dawn of the Virtual Revolution*. William Morrow, New York, NY.
- BLAUERT, J. 1996. *Spatial hearing: the psychophysics of human sound localization*. MIT Press, Cambridge, MA, USA.
- BOELLSTORFF, T. 2008. *Coming of age in Second Life: An anthropologist explores the virtually human*. Princeton University Press, Princeton, NJ.
- BONNEEL, N., DRETTAKIS, G., TSINGOS, N., VIAUD-DELMON, I., AND JAMES, D. 2008. Fast modal sounds with scalable frequency-domain synthesis. *ACM Trans. Graph.* 27, 24:1–24:9.
- BOULANGER, K., PATTANAIK, S. N., AND BOUATOUCH, K. 2009. Rendering grass in real time with dynamic lighting. *IEEE Comput. Graph. Appl.* 29, 32–41.
- BOUTHORS, A., NEYRET, F., MAX, N., BRUNETON, E., AND CRASSIN, C. 2008. Interactive multiple anisotropic scattering in clouds. In *Proceedings of the 2008 symposium on Interactive 3D graphics and games*. I3D '08. ACM, New York, NY, USA, 173–182.
- BRADLEY, D., HEIDRICH, W., POPA, T., AND SHEFFER, A. 2010. High resolution passive facial performance capture. *ACM Trans. Graph.* 29, 41:1–41:10.
- BROWN, C. P. AND DUDA, R. O. 1998. A structural model for binaural sound synthesis. *IEEE Transactions on Speech and Audio Processing* 6, 476–488.

- BULLION, C. AND GUROCAK, H. 2009. Haptic glove with mr brakes for distributed finger force feedback. *Presence: Teleoper. Virtual Environ.* 18, 421–433.
- BURNS, W. 2010. Defining the metaverse — revisited. <http://cityofnidus.blogspot.com/2010/04/defining-metaverse-revisited.html>. Accessed on May 17, 2011.
- CHADWICK, J. N., AN, S. S., AND JAMES, D. L. 2009. Harmonic shells: a practical nonlinear sound model for near-rigid thin shells. *ACM Trans. Graph.* 28, 119:1–119:10.
- CHANDRASIRI, N. P., NAEMURA, T., ISHIZUKA, M., HARASHIMA, H., AND BARAKONYI, I. 2004. Internet communication using real-time facial expression analysis and synthesis. *IEEE Multi-Media* 11, 20–29.
- CHARLES, J. 1999. Neural interfaces link the mind and the machine. *Computer* 32, 16–18.
- CHOW, Y.-W., POSE, R., REGAN, M., AND PHILLIPS, J. 2006. Human visual perception of region warping distortions. In *Proceedings of the 29th Australasian Computer Science Conference - Volume 48*. ACSC '06. Australian Computer Society, Inc., Darlinghurst, Australia, Australia, 217–226.
- CHUN, J., KWON, O., MIN, K., AND PARK, P. 2007. Real-time face pose tracking and facial expression synthesizing for the animation of 3d avatar. In *Proceedings of the 2nd international conference on Technologies for e-learning and digital entertainment*. Edutainment'07. Springer-Verlag, Berlin, Heidelberg, 191–201.
- CINQUETTI (AV), K. 2011. Kirstens Viewer web site. <http://www.kirstensviewer.com>. Accessed on February 17, 2011.
- COBOS, M., LOPEZ, J. J., AND SPORS, S. 2010. A sparsity-based approach to 3D binaural sound synthesis using time-frequency array processing. *EURASIP Journal on Advances in Signal Processing* 2010.
- COLLADA COMMUNITY. 2011. COLLADA: Digital Asset and FX Exchange Schema wiki. <http://collada.org>. Accessed on May 17, 2011.
- COZZI, L., D'ANGELO, P., CHIAPPALONE, M., IDE, A. N., NOVELLINO, A., MARTINOIA, S., AND SANGUINETI, V. 2005. Coding and decoding of information in a bi-directional neural interface. *Neurocomput.* 65-66, 783–792.
- CRUZ-NEIRA, C., SANDIN, D. J., DEFANTI, T. A., KENYON, R. V., AND HART, J. C. 1992. The cave: audio visual experience automatic virtual environment. *Commun. ACM* 35, 64–72.
- DE PASCALE, M., MULATTO, S., AND PRATTICHIZZO, D. 2008. Bringing haptics to Second Life. In *Proceedings of the 2008 Ambi-Sys workshop on Haptic user interfaces in ambient media systems*. HAS '08. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, 6:1–6:6.
- DEAN (AV), N. 2011. Nirans Viewer project site. <http://sourceforge.net/projects/niransviewer>. Accessed on December 15, 2011.
- ELEK, O. AND KMOCH, P. 2010. Real-time spectral scattering in large-scale natural participating media. In *Proceedings of the 26th Spring Conference on Computer Graphics*. SCCG '10. ACM, New York, NY, USA, 77–84.
- FACEBOOK. 2011. Facebook Developers web site. <http://developers.facebook.com>. Accessed on May 17, 2011.
- FISHER, B., FELS, S., MACLEAN, K., MUNZNER, T., AND RENSINK, R. 2004. Seeing, hearing, and touching: putting it all together. In *ACM SIGGRAPH 2004 Course Notes*. SIGGRAPH '04. ACM, New York, NY, USA.
- FOLGHERAITER, M., GINI, G., AND VERCESI, D. 2008. A multi-modal haptic interface for virtual reality and robotics. *J. Intell. Robotics Syst.* 52, 465–488.
- FRANCONE, J. AND NIGAY, L. 2011. 3D displays on mobile devices: head-coupled perspective. <http://iihm.imag.fr/en/demo/hcpmobile>. Accessed on April 19, 2011.
- FREY, D., ROYAN, J., PIEGAY, R., KERMARREC, A.-M., ANCEAUME, E., AND FESSANT, F. L. 2008. Solipsis: A decentralized architecture for virtual environments. In *1st International Workshop on Massively Multiuser Virtual Environments (MMVE)*. IEEE Workshop on Massively Multiuser Virtual Environments, Reno, NV, USA, 29–33.

- GAMASUTRA STAFF. 2011. Behind the scenes of trion worlds' rift. http://www.gamasutra.com/view/news/35009/Exclusive_Behind_The_Scenes_Of_Trion_Worlds_Rift.php. Accessed on June 24, 2011.
- GEARZ (AV), S. 2011. Singularity Viewer web site. <http://www.singularityviewer.org>. Accessed on March 30, 2011.
- GEDDES, L. 2010. Immortal avatars: Back up your brain, never die. *New Scientist* 206, 2763 (June), 28–31.
- GHOSH, A. 2007. Realistic materials and illumination environments. Ph.D. thesis, University of British Columbia, Vancouver, BC, Canada, Canada. AAINR31775.
- GILBERT, R. L. 2011. The P.R.O.S.E. Project: A program of in-world behavioral research on the Metaverse. *Journal of Virtual Worlds Research* 4, 1 (July), 3–18.
- GILBERT, R. L. AND FORNEY, A. 2011. The distributed self: Virtual worlds and the future of human identity. In *Postcards from the Metaverse — Reflections on how the virtual entangling with the physical may impact society, politics, and the economy*, R. Teigland and D. Power, Eds. In press.
- GILBERT, R. L., FOSS, J., AND MURPHY, N. 2011. Multiple personality order: Physical and personality characteristics of the self, primary avatar, and alt. In *Reinventing Ourselves: Contemporary Concepts of Identity in Virtual Worlds*, A. Peachey and M. Childs, Eds. Springer Publishing, London, 213–234.
- GILBERT, R. L., GONZALEZ, M. A., AND MURPHY, N. A. 2011. Sexuality in the 3D Internet and its relationship to real-life sexuality. *Psychology & Sexuality* 2, 107–122.
- GLADSTONE, D. 2011. Trion rift opens on HP servers. The Next Bench Blog, <http://h20435.www2.hp.com/t5/The-Next-Bench-Blog/Trion-Rift-Opens-on-HP-Servers/ba-p/61333>. Accessed on June 24, 2011.
- GLENCROSS, M., CHALMERS, A. G., LIN, M. C., OTADUY, M. A., AND GUTIERREZ, D. 2006. Exploiting perception in high-fidelity virtual environments. In *ACM SIGGRAPH 2006 Courses*. SIGGRAPH '06. ACM, New York, NY, USA.
- GOSSELIN, B. AND SAWAN, M. 2010. A low-power integrated neural interface with digital spike detection and extraction. *Analog Integr. Circuits Signal Process.* 64, 3–11.
- GREENBERG, S., MARQUARDT, N., BALLENDAT, T., DIAZ-MARINO, R., AND WANG, M. 2011. Proxemic interactions: the new ubicomp? *interactions* 18, 42–50.
- GUPTA, N., DEMERS, A., GEHRKE, J., UNTERBRUNNER, P., AND WHITE, W. 2009. Scalability for virtual worlds. In *Data Engineering, 2009. ICDE '09. IEEE 25th International Conference on*. 1311–1314.
- HAANS, A. AND IJSSELSTEIJN, W. 2006. Mediated social touch: a review of current research and future directions. *Virtual Real.* 9, 149–159.
- HAO, A., SONG, F., LI, S., LIU, X., AND XU, X. 2009. Real-time realistic rendering of tissue surface with mucous layer. In *Proceedings of the 2009 Second International Workshop on Computer Science and Engineering - Volume 01. IWCSE '09*. IEEE Computer Society, Washington, DC, USA, 302–306.
- HARDERS, M., BIANCHI, G., AND KNOERLEIN, B. 2007. Multimodal augmented reality in medicine. In *Proceedings of the 4th international conference on Universal access in human-computer interaction: ambient interaction*. UAHCI'07. Springer-Verlag, Berlin, Heidelberg, 652–658.
- HAYHOE, M. M., BALLARD, D. H., TRIESCH, J., SHINODA, H., AIVAR, P., AND SULLIVAN, B. 2002. Vision in natural and virtual environments. In *Proceedings of the 2002 symposium on Eye tracking research & applications*. ETRA '02. ACM, New York, NY, USA, 7–13.
- HERTZMANN, A., O'SULLIVAN, C., AND PERLIN, K. 2009. Realistic human body movement for emotional expressiveness. In *ACM SIGGRAPH 2009 Courses*. SIGGRAPH '09. ACM, New York, NY, USA, 20:1–20:27.
- HO, C.-C., MACDORMAN, K. F., AND PRAMONO, Z. A. D. D. 2008. Human emotion and the uncanny valley: a glm, mds, and isomap analysis of robot video ratings. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*. HRI '08. ACM, New York, NY, USA, 169–176.

- HU, Y., VELHO, L., TONG, X., GUO, B., AND SHUM, H. 2006. Realistic, real-time rendering of ocean waves: Research articles. *Comput. Animat. Virtual Worlds* 17, 59–67.
- IEEE VW STANDARD WORKING GROUP. 2011a. IEEE Virtual World Standard Working Group wiki. <http://www.metaversestandards.org>. Accessed on May 17, 2011.
- IEEE VW STANDARD WORKING GROUP. 2011b. Terminology and definitions. http://www.metaversestandards.org/index.php?title=Terminology_and_Definitions#MetaWorld. Accessed on May 17, 2011.
- IMAGE METRICS. 2011. Image Metrics web site. <http://www.image-metrics.com>. Accessed on April 20, 2011.
- JAY, C., GLENCROSS, M., AND HUBBOLD, R. 2007. Modeling the effects of delayed haptic and visual feedback in a collaborative virtual environment. *ACM Trans. Comput.-Hum. Interact.* 14, 8/1–8/31.
- JEON, N. J., LEEM, C. S., KIM, M. H., AND SHIN, H. G. 2007. A taxonomy of ubiquitous computing applications. *Wirel. Pers. Commun.* 43, 1229–1239.
- JIMENEZ, J., SCULLY, T., BARBOSA, N., DONNER, C., ALVAREZ, X., VIEIRA, T., MATTS, P., ORVALHO, V., GUTIERREZ, D., AND WEYRICH, T. 2010. A practical appearance model for dynamic facial color. *ACM Trans. Graph.* 29, 141:1–141:10.
- JOLY, J. 2010. Can a polygon make you cry? <http://www.jonathanjoly.com/front.htm>. Accessed on May 17, 2011.
- KALL BINAURAL AUDIO. 2010. A brief history of binaural audio. <http://www.kallbinauralaudio.com/binauralhistory.html>. Accessed on March 7, 2011.
- KAWATO, S. AND OHYA, J. 2000. Real-time detection of nodding and head-shaking by directly detecting and tracking the “between-eyes”. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition 2000*. FG ’00. IEEE Computer Society, Washington, DC, USA, 40–.
- KELLER, J. AND SIMON, G. 2002. Toward a peer-to-peer shared virtual reality. In *Proceedings of the 22nd International Conference on Distributed Computing Systems*. ICDCSW ’02. IEEE Computer Society, Washington, DC, USA, 695–700.
- KIM, J.-S., GRAČANIN, D., MATKOVIĆ, K., AND QUEK, F. 2010. The effects of finger-walking in place (FWIP) for spatial knowledge acquisition in virtual environments. In *Proceedings of the 10th international conference on Smart graphics*. SG’10. Springer-Verlag, Berlin, Heidelberg, 56–67.
- KOEPNICK, S., HOANG, R. V., SGAMBATI, M. R., COMING, D. S., SUMA, E. A., AND SHERMAN, W. R. 2010. Graphics for serious games: Rist: Radiological immersive survey training for two simultaneous users. *Comput. Graph.* 34, 665–676.
- KOROLOV, M. 2011. Kitley brings Facebook, instant regions to OpenSim. *Hypergrid Business*.
- KROL, L. R., ALIAKSEYEU, D., AND SUBRAMANIAN, S. 2009. Haptic feedback in remote pointing. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*. CHI EA ’09. ACM, New York, NY, USA, 3763–3768.
- KRUEGER, M. W. 1993. *An Easy Entry Artificial Reality*. Academic Press Professional, Cambridge, MA, Chapter 7, 147–162.
- KUCHENBECKER, K. J. 2008. Haptography: capturing the feel of real objects to enable authentic haptic rendering (invited paper). In *Proceedings of the 2008 Ambi-Sys workshop on Haptic user interfaces in ambient media systems*. HAS ’08. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, 3:1–3:3.
- KUMAR, S., CHHUGANI, J., KIM, C., KIM, D., NGUYEN, A., DUBEY, P., BIENIA, C., AND KIM, Y. 2008. Second life and the new generation of virtual worlds. *Computer* 41, 46–53.
- KURZWEIL, R. 2001. The law of accelerating returns. <http://www.kurzweilai.net/the-law-of-accelerating-returns>. Accessed on April 11, 2011.
- LANIER, J. 1992. Virtual reality: the promise of the future. *Interact. Learn. Int.* 8, 275–279.
- LANIER, J. 2010. *You are not a Gadget: A Manifesto*. Alfred A. Knopf, New York, NY, USA.
- LANIER, J. AND BIOCICA, F. 1992. An insider’s view of the future of virtual reality. *Journal of Communication* 42, 4 (Autumn), 150–172.

- LATOSCHIK, M. E. 2005. A user interface framework for multimodal VR interactions. In *Proceedings of the 7th international conference on Multimodal interfaces*. ICMI '05. ACM, New York, NY, USA, 76–83.
- LÉCUYER, A. 2009. Simulating haptic feedback using vision: A survey of research and applications of pseudo-haptic feedback. *Presence: Teleoper. Virtual Environ.* 18, 39–53.
- LEE, Y., WAMPLER, K., BERNSTEIN, G., POPOVIĆ, J., AND POPOVIĆ, Z. 2010. Motion fields for interactive character locomotion. *ACM Trans. Graph.* 29, 138:1–138:8.
- LEWINSKI, J. S. 2011. Light at the end of the racetrack: How Pixar explored the physics of light for *Cars 2*. *Scientific American*, 20.
- LINDEN LAB. 2011a. Extended second life knowledge base: Graphics preferences layout. http://wiki.secondlife.com/wiki/Graphics_Preferences_Layout. Accessed on April 5, 2011.
- LINDEN LAB. 2011b. New sl viewer with improved search and real-time shadows. <http://community.secondlife.com/t5/Featured-News/New-SL-Viewer-with-Improved-Search-and-Real-Time-Shadows/ba-p/927463>. Accessed on July 12, 2011.
- LIU, H., BOWMAN, M., ADAMS, R., HURLIMAN, J., AND LAKE, D. 2010. Scaling virtual worlds: Simulation requirements and challenges. In *Winter Simulation Conference (WSC), Proceedings of the 2010*. The WSC Foundation, Baltimore, MD, USA, 778–790.
- LOBER, A., SCHWABE, G., AND GRIMM, S. 2007. Audio vs. chat: The effects of group size on media choice. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*. HICSS '07. IEEE Computer Society, Washington, DC, USA, 41–.
- LUDLOW, P. AND WALLACE, M. 2007. *The Second Life Herald: The virtual tabloid that witnessed the dawn of the Metaverse*. MIT Press, New York.
- LYNN, R. 2004. Ins and outs of teledildonics. *Wired*. Accessed on March 30, 2011.
- LYON (AV), J. 2011. Phoenix Viewer web site. <http://www.phoenixviewer.com>. Accessed on March 30, 2011.
- MARCOS, S., GÓMEZ-GARCÍA-BERMEJO, J., AND ZALAMA, E. 2010. A realistic, virtual head for human-computer interaction. *Interact. Comput.* 22, 176–192.
- MARTIN, J.-C., D'ALESSANDRO, C., JACQUEMIN, C., KATZ, B., MAX, A., POINTAL, L., AND RILLIARD, A. 2007. 3d audiovisual rendering and real-time interactive control of expressivity in a talking head. In *Proceedings of the 7th international conference on Intelligent Virtual Agents*. IVA '07. Springer-Verlag, Berlin, Heidelberg, 29–36.
- MCDONNELL, R. AND BREIDT, M. 2010. Face reality: investigating the uncanny valley for virtual faces. In *ACM SIGGRAPH ASIA 2010 Sketches*. SA '10. ACM, New York, NY, USA, 41:1–41:2.
- MCDONNELL, R., JÖRG, S., MCHUGH, J., NEWELL, F. N., AND O'SULLIVAN, C. 2009. Investigating the role of body shape on the perception of emotion. *ACM Trans. Appl. Percept.* 6, 14:1–14:11.
- MECWAN, A. I., SAVANI, V. G., RAJVI, S., AND VAYA, P. 2009. Voice conversion algorithm. In *Proceedings of the International Conference on Advances in Computing, Communication and Control*. ICAC3 '09. ACM, New York, NY, USA, 615–619.
- MISSELHORN, C. 2009. Empathy with inanimate objects and the uncanny valley. *Minds Mach.* 19, 345–359.
- MOORE, G. E. 1965. Cramming more components into integrated circuits. *Electronics* 38, 8 (April).
- MORI, M. 1970. Bukimi no tani (the uncanny valley). *Energy* 7, 4, 33–35.
- MORNINGSTAR, C. AND FARMER, F. R. 1991. *The lessons of Lucasfilm's habitat*. MIT Press, Cambridge, MA, USA, 273–302.
- MORRIS, D., KELLAND, M., AND LLOYD, D. 2005. *Machinima: Making Animated Movies in 3D Virtual Environments*. Muska & Lipman/Premier-Trade, Cincinnati, OH.
- MOSS, W., YEH, H., HONG, J.-M., LIN, M. C., AND MANOCHA, D. 2010. Sounding liquids: Automatic sound synthesis from fluid simulation. *ACM Trans. Graph.* 29, 21:1–21:13.
- NAKAMOTO, S. 2009. Bitcoin: A peer-to-peer electronic cash system. <http://www.bitcoin.org/bitcoin.pdf>. Accessed on June 26, 2011.

- NATIONAL ACADEMY OF ENGINEERING. 2008. *Grand Challenges for Engineering*. National Academy of Engineering, Washington, DC, USA.
- NAUMANN, A. B., WECHSUNG, I., AND HURTIENNE, J. 2010. Multimodal interaction: A suitable strategy for including older users? *Interact. Comput.* 22, 465–474.
- NEWMAN, R., MATSUMOTO, Y., ROUGEAX, S., AND ZELINSKY, A. 2000. Real-time stereo tracking for head pose and gaze estimation. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition 2000*. FG '00. IEEE Computer Society, Washington, DC, USA, 122–.
- OAuth COMMUNITY. 2011. OAuth home page. <http://oauth.net>. Accessed on May 17, 2011.
- O'BRIEN, J. F., SHEN, C., AND GATCHALIAN, C. M. 2002. Synthesizing sounds from rigid-body simulations. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*. SCA '02. ACM, New York, NY, USA, 175–181.
- OLANO, M., AKELEY, K., HART, J. C., HEIDRICH, W., MCCOOL, M., MITCHELL, J. L., AND ROST, R. 2004. Real-time shading. In *ACM SIGGRAPH 2004 Course Notes*. SIGGRAPH '04. ACM, New York, NY, USA.
- OLANO, M. AND LASTRA, A. 1998. A shading language on graphics hardware: the pixelflow shading system. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. SIGGRAPH '98. ACM, New York, NY, USA, 159–168.
- OPEN COBALT PROJECT. 2011. Open Cobalt web site. <http://www.opencobalt.org>. Accessed on March 30, 2011.
- OPEN WONDERLAND FOUNDATION. 2011. Open Wonderland web site. <http://www.openwonderland.org>. Accessed on March 30, 2011.
- OPENID FOUNDATION. 2011. OpenID Foundation home page. <http://openid.net>.
- OPENSIMULATOR PROJECT. 2011. OpenSimulator web site. <http://www.opensimulator.org>. Accessed on March 30, 2011.
- OTADUY, M. A., IGARASHI, T., AND LAVIOLA, JR., J. J. 2009. Interaction: interfaces, algorithms, and applications. In *ACM SIGGRAPH 2009 Courses*. SIGGRAPH '09. ACM, New York, NY, USA, 14:1–14:66.
- PALEARI, M. AND LISETTI, C. L. 2006. Toward multimodal fusion of affective cues. In *Proceedings of the 1st ACM international workshop on Human-centered multimedia*. HCM '06. ACM, New York, NY, USA, 99–108.
- PIOTROWSKI, J. AND TERZIS, G. 2011. Virtual currency: Bits and bob. <http://www.economist.com/blogs/babbage/2011/06/virtual-currency>. Accessed on June 26, 2011.
- POCKET METAVERSE. 2009. Pocket Metaverse web site. <http://www.pocketmetaverse.com>. Accessed on May 30, 2011.
- RAHMAN, A. S. M. M., HOSSAIN, S. K. A., AND SADDIK, A. E. 2010. Bridging the gap between virtual and real world by bringing an interpersonal haptic communication system in Second Life. In *Proceedings of the 2010 IEEE International Symposium on Multimedia*. ISM '10. IEEE Computer Society, Washington, DC, USA, 228–235.
- RANKINE, R., AYROMLOU, M., ALIZADEH, H., AND RIVERS, W. 2011. Kinect Second Life interface. http://www.youtube.com/watch?v=KvVWjah_fXU. Accessed on May 31, 2011.
- RAYMOND, E. S. 2001. *The cathedral and the bazaar: musings on Linux and open source by an accidental revolutionary*. O'Reilly & Associates, Inc., Sebastopol, CA, USA.
- REALXTEND ASSOCIATION. 2011. realXtend web site. <http://realxtend.wordpress.com>. Accessed on December 15, 2011.
- RECANZONE, G. H. 2003. Auditory influences on visual temporal rate perception. *Journal of Neurophysiology* 89, 2, 1078–1093.
- REDDWARF SERVER PROJECT. 2011. RedDwarf Server web site. <http://www.reddwarfserver.org>. Accessed on June 14, 2011.
- REED, D. P. 2005. Designing croquet's teatime: a real-time, temporal environment for active object cooperation. In *Companion to the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*. OOPSLA '05. ACM, New York, NY, USA, 7–7.

- ROH, C. H. AND LEE, W. B. 2005. Image blending for virtual environment construction based on tip model. In *Proceedings of the 4th WSEAS International Conference on Signal Processing, Robotics and Automation*. World Scientific and Engineering Academy and Society (WSEAS), Stevens Point, Wisconsin, USA, 26:1–26:5.
- ROST, R. J., LICEA-KANE, B., GINSBURG, D., KESSENICH, J. M., LICHTENBELT, B., MALAN, H., AND WEIBLEN, M. 2009. *OpenGL Shading Language*, 3rd ed. Addison-Wesley Professional, Boston, MA.
- SALLNÄS, E.-L. 2010. Haptic feedback increases perceived social presence. In *Proceedings of the 2010 international conference on Haptics - generating and perceiving tangible sensations: Part II*. EuroHaptics'10. Springer-Verlag, Berlin, Heidelberg, 178–185.
- SCHMIDT, K. O. AND DIONISIO, J. D. N. 2008. Texture map based sound synthesis in rigid-body simulations. In *ACM SIGGRAPH 2008 posters*. ACM, New York, NY, USA, 11:1–11:1.
- SENG, D. AND WANG, H. 2009. Realistic real-time rendering of 3d terrain scenes based on opengl. In *Proceedings of the 2009 First IEEE International Conference on Information Science and Engineering*. ICISE '09. IEEE Computer Society, Washington, DC, USA, 2121–2124.
- SEYAMA, J. AND NAGAYAMA, R. S. 2007. The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoper. Virtual Environ.* 16, 337–351.
- SEYAMA, J. AND NAGAYAMA, R. S. 2009. Probing the uncanny valley with the eye size aftereffect. *Presence: Teleoper. Virtual Environ.* 18, 321–339.
- SHAH, M. 2007. Image-space approach to real-time realistic rendering. Ph.D. thesis, University of Central Florida, Orlando, FL, USA. AAI3302925.
- SHAH, M. A., KONTTINEN, J., AND PATTANAIK, S. 2009. Image-space subsurface scattering for interactive rendering of deformable translucent objects. *IEEE Comput. Graph. Appl.* 29, 66–78.
- SHORT, J., WILLIAMS, E., AND CHRISTIE, B. 1976. *The social psychology of telecommunications*. Wiley, London.
- SMART, J., CASCIO, J., AND PATTENDORF, J. 2007. Metaverse roadmap overview. <http://www.metaverseroadmap.org>. Accessed on February 16, 2011.
- SMURRAYINCHESTER (AV) AND VOIDVECTOR (AV). 2008. Mori uncanny valley.svg. http://commons.wikimedia.org/wiki/File:Mori_Uncanny_Valley.svg under the Creative Commons Attribution-Share Alike 3.0 license. Accessed on June 27, 2011.
- STOLL, C., GALL, J., DE AGUIAR, E., THRUN, S., AND THEOBALT, C. 2010. Video-based reconstruction of animatable human characters. *ACM Trans. Graph.* 29, 139:1–139:10.
- SUBRAMANIAN, S., GUTWIN, C., NACENTA, M., POWER, C., AND JUN, L. 2005. Haptic and tactile feedback in directed movements. In *Proceedings of the Conference on Guidelines on Tactile and Haptic Interactions*. Saskatoon, Saskatchewan, Canada.
- SUN, W. AND MUKHERJEE, A. 2006. Generalized wavelet product integral for rendering dynamic glossy objects. *ACM Trans. Graph.* 25, 955–966.
- SUNG, J. AND KIM, D. 2009. Real-time facial expression recognition using STAAM and layered GDA classifier. *Image Vision Comput.* 27, 1313–1325.
- TAYLOR, M. T., CHANDAK, A., ANTANI, L., AND MANOCHA, D. 2009. RESound: interactive sound rendering for dynamic virtual environments. In *Proceedings of the 17th ACM international conference on Multimedia*. MM '09. ACM, New York, NY, USA, 271–280.
- TINWELL, A. AND GRIMSHAW, M. 2009. Bridging the uncanny: an impossible traverse? In *Proceedings of the 13th International MindTrek Conference: Everyday Life in the Ubiquitous Era*. MindTrek '09. ACM, New York, NY, USA, 66–73.
- TINWELL, A., GRIMSHAW, M., NABI, D. A., AND WILLIAMS, A. 2011. Facial expression of emotion and perception of the uncanny valley in virtual characters. *Comput. Hum. Behav.* 27, 741–749.
- TRATTNER, C., STEURER, M. E., AND KAPPE, F. 2010. Socializing virtual worlds with facebook: a prototypical implementation of an expansion pack to communicate between facebook and open-simulator based virtual worlds. In *Proceedings of the 14th International Academic MindTrek Conference: Envisioning Future Media Environments*. MindTrek '10. ACM, New York, NY, USA, 158–160.

- TRIESCH, J. AND VON DER MALSBERG, C. 1998. Robotic gesture recognition by cue combination. In *Informatik '98: Informatik zwischen Bild und Sprache*, J. Dassow and R. Kruse, Eds. Springer-Verlag, Magdeburg, Germany.
- TSETSERUKOU, D. 2010. HaptiHug: a novel haptic display for communication of hug over a distance. In *Proceedings of the 2010 international conference on Haptics: generating and perceiving tangible sensations, Part I*. EuroHaptics'10. Springer-Verlag, Berlin, Heidelberg, 340–347.
- TSETSERUKOU, D., NEVIAROUSKAYA, A., PRENDINGER, H., ISHIZUKA, M., AND TACHI, S. 2010. iFeel_IM: innovative real-time communication system with rich emotional and haptic channels. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*. CHI EA '10. ACM, New York, NY, USA, 3031–3036.
- TURKLE, S. 1995. *Life on the Screen: Identity in the Age of the Internet*. Simon & Schuster Trade, New York, NY, USA.
- TURKLE, S. 2007. *Evocative Objects: Things We Think With*. MIT Press, Cambridge, MA, USA.
- TWITTER. 2011. Twitter API Documentation web site. <http://dev.twitter.com/doc>. Accessed on May 17, 2011.
- UEDA, Y., HIROTA, M., AND SAKATA, T. 2006. Vowel synthesis based on the spectral morphing and its application to speaker conversion. In *Proceedings of the First International Conference on Innovative Computing, Information and Control - Volume 2*. ICICIC '06. IEEE Computer Society, Washington, DC, USA, 738–741.
- VAN DEN DOEL, K., KRY, P. G., AND PAI, D. K. 2001. FoleyAutomatic: physically-based sound effects for interactive simulation and animation. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. SIGGRAPH '01. ACM, New York, NY, USA, 537–544.
- VIGNAIS, N., KULPA, R., CRAIG, C., BRAULT, S., MULTON, F., AND BIDEAU, B. 2010. Influence of the graphical levels of detail of a virtual thrower on the perception of the movement. *Presence: Teleoper. Virtual Environ.* 19, 243–252.
- WACHS, J. P., KÖLSCH, M., STERN, H., AND EDAN, Y. 2011. Vision-based hand-gesture applications. *Commun. ACM* 54, 60–71.
- WADLEY, G., GIBBS, M. R., AND DUCHENEAUT, N. 2009. You can be too rich: mediated communication in a virtual world. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7*. OZCHI '09. ACM, New York, NY, USA, 49–56.
- WALDO, J. 2008. Scaling in games and virtual worlds. *Commun. ACM* 51, 38–44.
- WANG, L., LIN, Z., FANG, T., YANG, X., YU, X., AND KANG, S. B. 2006. Real-time rendering of realistic rain. In *ACM SIGGRAPH 2006 Sketches*. SIGGRAPH '06. ACM, New York, NY, USA.
- WEISER, M. 1991. The computer for the 21st century. *Scientific American* 265, 94–102.
- WEISER, M. 1993. Ubiquitous computing. *Computer* 26, 71–72.
- WILKES, I. 2008. Second life: How it works (and how it doesn't). QCon 2008, <http://www.infoq.com/presentations/Second-Life-Ian-Wilkes>. Accessed on June 13, 2011.
- WITHANA, A., KONDO, M., MAKINO, Y., KAKEHI, G., SUGIMOTO, M., AND INAMI, M. 2010. ImpAct: Immersive haptic stylus to enable direct touch and manipulation for surface computing. *Comput. Entertain.* 8, 9:1–9:16.
- XU, N., SHAO, X., AND YANG, Z. 2008. A novel voice morphing system using bi-gmm for high quality transformation. In *Proceedings of the 2008 Ninth ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing*. IEEE Computer Society, Washington, DC, USA, 485–489.
- YAHOO! INC. 2011. Flickr Services web site. <http://www.flickr.com/services/api>. Accessed on May 17, 2011.
- ZHAO, Q. 2011. 10 scientific problems in virtual reality. *Commun. ACM* 54, 116–118.
- ZHENG, C. AND JAMES, D. L. 2009. Harmonic fluids. *ACM Trans. Graph.* 28, 37:1–37:12.
- ZHENG, C. AND JAMES, D. L. 2010. Rigid-body fracture sound with precomputed soundbanks. *ACM Trans. Graph.* 29, 69:1–69:13.

